An Economic Model of Prior-Free Spatial Search

Suraj Malladi*

April 22, 2022

Abstract

We propose a model of sequential spatial search with learning. There is a mapping from a space of technologies (or products) to qualities that is unknown to the searcher. The searcher can learn various points on this mapping through costly experimentation. She cares both about the technology that she adopts as well as the best one available, as would a firm in an innovation race or an online shopper concerned with missing a good deal. She does not have a prior over mappings but knows only that neighboring technologies in attribute space are similar in quality. We characterize optimal search strategies when the searcher worries about worst-case mappings at every step of the way. These are mappings that trigger wild-goose chases: excessive search with relatively poor discoveries to show for it. We derive comparative statics that match patterns observed in empirical studies on spatial search. Finally, we apply the results to the problem of optimal search space design faced by online platforms.

*Thanks to Tommaso Denti, David Easley and Andrzej Skrzypacz for their feedback and suggestions. I am particularly grateful to Ilya Morozov for so many helpful conversations. Cornell University. Email: surajm@cornell.edu.

1 Introduction

People search for various things: ideas for a startup, ideas for a paper, a new set of golf clubs. They examine some items in detail and learn how good they are. In the process, they also get a sense for the quality of similar, unexplored alternatives. These searchers dynamically reassess which alternatives to consider next or whether to settle for their best discoveries.

There are no flexible and tractable models of such behavior with rational agents. Classic models of costly sequential search either treat search as a pure stopping problem (McCall, 1970; Rothschild, 1978) or assume that search outcomes are independent, shutting down scope for learning (Weitzman, 1979). Models where agents learn and choose the order of search are limited to the case where payoffs for different objects are independent conditional on unknown parameters (Adam, 2001). Models with richer *spatial learning*, where values of proximate objects in attribute space are more correlated, exogenously fix the order of search (Urgun and Yariv, 2021; Wong, 2021). Models with spatial learning and endogenous search order address the case where agents have limited look-ahead (Callander, 2011; Garfagnini and Strulovici, 2016; Hodgson and Lewis, 2020).

Missing in the search theory literature is a tractable model with spatial learning where fully forward-looking agents choose their order of search. Such a model is needed to (1) derive comparative statics, (2) study optimal policies aimed at designing a search space, (3) estimate search costs from spatial search data, under the standard economics assumption "that individuals are maximizing something, even if the something is unorthodox" (Lazear, 2000). This paper puts forth one such model and derives comparative statics that tie out empirical findings in Bronnenberg et al. (2016), Blake et al. (2016), and Hodgson and Lewis (2020). The following examples motivate the exercise:

Firm R&D A researcher at a startup experiments with the design of a novel technology. She prototypes different configurations in attribute space at different costs. She knows that small tweaks to the earlier designs can only affect quality by so much, but otherwise entertains many possibilities for how the true mapping from attributes to design quality looks. The success of the startup is predicated both on its own design quality and that of an incumbent who is privately searching in the same space. The incumbent has many resources to spend on R&D, so the startup expects to compete against a close to ideal design.

Dose-ranging study A pharmaceutical company is developing a treatment involving a new compound and again is concerned with a competitor. In addition to knowing that similar dosages have similar effects, the mapping from dosage to treatment quality is thought to have a sweet spot. Too high a dose leads to toxicity, too low a dose is ineffective; the best outcomes are near the middle of some 'therapeutic window'. What is too high or low and whether the compound is even effective or toxic at all is determined through costly experimentation.

Consumer search A consumer browses a list of TVs ranked by price online, hoping to find a bargain. He expects similarly priced TVs to be of similar quality and can learn his expected surplus for a product by clicking and reading about it. He is happy to make a good purchase but dislikes learning afterward that he missed out on a much better deal. He searches and expends costly attention to mitigate this scenario.

The base model considers a searcher who searches (or experiments with) items (or technologies) in some compact, one-dimensional attribute space. The searcher does not know the true mapping from attributes to quality, i.e., the quality index. She knows only that it belongs to some set of possible mappings. She sequentially experiments with different technologies to learn their true quality, and in the process, narrows down the set of rationalizable quality indices. After each experiment, the searcher decides whether to continue searching or settle with the best technology she had discovered so far.

The searcher's anticipated payoff increases in the quality of the best technology she discovers, decreases in the quality of the best attainable technology, and decreases with costly search. At each turn, she makes a decision to maximize her worst-case anticipated payoff. She takes this worst-case over the rationalizable quality indices at that point in time. In effect, she fears going on a wild goose chase only to miss an unexplored high quality technology. But the more she learns, the less scope there is for this unhappy outcome: if a nearby option were so good, then what she had found could not have been too far off.

In analyzing the model, we focus on two forms of prior information.

In the case of *minimal prior knowledge*, the searcher entertains the possibility that the quality index could be any Lipschitz continuous mapping for some Lipschitz constant. This captures the spatial element of learning by bounding how much the quality of an undiscovered technology can vary, given what the searcher knows about nearby discovered technologies. The Lipschitz constant represents the searcher's perception of search complexity. For example, a startup may perceive high complexity if it is developing an application where small tweaks to user interface can trigger drastic user responses.

The second case is that the searcher entertains the possibility that the quality index could be any Lipschitz continuous and *quasiconcave* mapping. This is to capture the drug development example where there may be a notion of a correct range of dosage levels, and moving away from this range leads to a deterioration in treatment quality. Complexity has the same interpretation as before. For example, the compound lithium is thought to have a particularly narrow therapeutic window, so discovering the right dosage for a new treatment can be like finding a needle in a haystack.

The first set of results characterize optimal search procedures when the searcher has minimal prior information or additionally knows quality indices are quasiconcave. Consider for a moment an alternative framework, akin to the classical model of simultaneous search

(Stigler, 1961), where the searcher must irrevocably decide the sequence of searches she will make beforehand and pick the best of the discovered technologies. Solving this alternative model to find an optimal simultaneous search procedure is a simpler task, because it rules out learning. We show that an optimal sequential search procedure can be found as follows: in each period, solve for the optimal simultaneous search procedure given the history of observations, and then take only the first step prescribed by the simultaneous procedure (rather than follow the full plan). Moreover, all optimal sequential search procedures can be found this way.

To see why, note that if the searcher commits to the sequence of technologies to explore, the worst-case outcome simply keeps the returns to search low and the value of the best unexplored alternative high. In a sequential search setting, the adversary additionally tries to limit how much the searcher learns to make her expend more wasteful effort searching. The result follows because the strategy that limits payoffs to search is precisely the strategy that limits searcher information.

For a simple intuition, suppose that the searcher already knows the quality of one technology. The worst-case outcome, ignoring search costs, is that all the options she explores are of this same quality. This would ensure that the best possible unexplored technology is as high quality as possible, conditional on the searcher failing to discover anything better than what she started with. Incidentally, this outcome also leaves the searcher the fewest clues for where to search next. By contrast, discovering a technology of very low quality would indicate to the searcher that she can avoid looking at nearby products; discovering a better technology may keep her search efforts in that area.

This characterization of optimal search allows for studying various comparative statics.

One basic comparative static is that the propensity to search is non-monotone in search complexity. Intuitively, if search complexity is low, then quality indices must be relatively flat. So if the searcher starts out with some technology, there is little value in paying the cost to experiment and innovate. On the other hand, if complexity is very high, then in the worst case, the searcher could experiment very aggressively and still miss out on good opportunities. Such a searcher is discouraged from engaging in search and going on a wild goose chase. If there are decreasing returns to being near the peak of the quality index, then search only takes place in some intermediate region.

We analyze how limiting look-ahead affects search dynamics. A myopic searcher, who lives only two periods, concludes search whenever an optimal forward-looking searcher would do the same. Moreover, a myopic searcher always searches in the most promising area of the search space. This is the region where a rationalizable quality index is as large as possible (i.e., the most unexplored area). When search costs are constant, a myopic searcher behaves like an optimistic searcher who believes that each search will yield the best-case outcome, rather than worst-case. Both engage in more radical experimentation than a forward-looking searcher, one to cover her bases and the other expecting to strike gold.

Next we show conditions under which experiments with sufficiently bad outcomes cause the searcher to either end experimentation or take a larger step size in technology space than in the previous period. This ties out an empirical observation by Hodgson and Lewis (2020) that there is a bigger jump away from unpopular products in online consumer search data. Suppose searchers are forward-looking but are incorrectly modeled as being myopic. An outside observer may misattribute a large jump in search space to myopia, underestimating how bad the last realized outcome was for the searcher.

The impact of news is more striking in the case where the searcher knows the quality index is quasiconcave. In this case, any experiment that reveals a better quality than hitherto observed stops the searcher from ever looking beyond the previous experiment. And if the searcher continues experimenting and the next experiment is worse, the searcher changes direction. Taken together, this produces the funnel-like search dynamics empirically observed by Bronnenberg et al. (2016) and Blake et al. (2016).

The baseline one-dimensional search space model is useful to develop intuitions about search dynamics in a prior free model. However, many search spaces are naturally multidimensional. For example, a shopper on an online marketplace may readily see the resolution, size, price and brand of different TVs even before he clicks on the product to learn more.

There are two natural ways to generalize the model when the search space is multidimensional. Sometimes, the searcher may learn her values for the attributes that make up an item. Other times, she may learn only her value for the item as a whole, without understanding the contribution of the constituent parts. We extend results along both directions.

We give an explicit algorithm for optimal multidimensional search when there is a finite set of objects search costs are constant. About optimal search with correlated rewards, Weitzman conjectures that:

> It appears plausible that other things being equal it would be better to open a box whose reward is highly correlated with other rewards because this adds a positive informational externality. But translating such an effect into a simple search rule seems difficult except in the most elementary cases. (Weitzman, 1979)

Our algorithm formalizes a prior-free analogue of Weitzman's conjecture. The searcher picks her targets to minimize a modified measure of distance between the set of explored and unexplored items. Proximity in attribute space roughly corresponds to higher correlation in a Bayesian model.

The algorithm takes polynomial many steps in the number of items in the search space, keeping search costs fixed. One may question the plausibility of models where agents optimally solve NP-hard problems, but our prior-free searcher avoids such troubles.

Finally, we consider the problem of an online marketplace designer who wants to order 'related products' to facilitate multidimensional search. We demonstrate the applicability of the model in a version of the problem where the designer wants to ensure that searchers of any type can follow their worst-case optimal search paths. The designer can sometimes solve this problem without any information about the searcher's preferences or perception of search complexity.

1.1 Related Literature

This paper belongs to the literature on search with learning.

One of the earliest papers on the topic is Rothschild (1978), which considers a model where an agent draws independent samples from an unknown distribution and decides when to stop and take the most recent draw. The agent has a multinomial Dirichlet prior about this distribution, which she updates over time. Bikhchandani and Sharma (1996) and the references therein generalize these results to allow for recall and other distributions. Schlag and Zapechelnyuk (2021) take prior-free approach to sequential search, motivated by the fact that solutions are only known in special cases and sensitive to distributional assumptions.

Introducing learning is more challenging in a model of search where items are heterogeneous ex-ante and agents with recall choose order of search. Weitzman (1979) applies the index solution for multi-arm bandits in Gittins (1974) to solve such a model without learning. Adam (2001) generalizes Rothschild (1978) and Weitzman (1979) to allow for multiple types of items. The main restriction is that all items have independently drawn payoffs. Learning about an item is informative about other items of the same type as their payoffs are drawn from the same unknown distribution.

More recently, papers have studied search and learning in environments with rich correlation structures. Notably, Callander (2011) models the mapping from the search space to outcomes as the realization of a Brownian motion path, linking strength of correlation to distance in search space. Short-lived agents decide where to search next, learning from the searches of their predecessors. Wong (2021) and Urgun and Yariv (2021) consider search on a Brownian motion by forward looking agents in settings where the searcher moves continuously in a fixed direction and chooses speed or scope of search.

A key feature of Callander (2011) (and Weitzman (1979)) is that agents are free to explore any point in the search space at any time. But incorporating this feature in a model of optimal search with learning has proven to be challenging. Garfagnini and Strulovici (2016) extend Callander (2011) to the case where agents search for two periods. But solving Callander (2011) with forward-looking searcher remains "an important open problem" (Hörner and Skrzypacz, 2017). The present paper solves a prior-free analogue of this problem.

A growing body of work in economics uses maximin or minimax regret objectives as a way to explain or advocate for simple and intuitive policies. This approach has been fruitful in topics like contracting (Carroll, 2019), regulation (Guo and Shmaya, 2019; Malladi, 2021), treatment choice (Manski, 2007; Tetenov, 2012), portfolio choice (Chassang, 2016) and experimental design Banerjee et al. (2017). We take a prior-free approach to optimal search. Bayesian learning is computationally intractable in search over correlated items, a drawback for a positive model. When search costs are constant, optimal search in our setting is computationally feasible and formalizes Weitzman's intuition that it is optimal to explore those objects which are most correlated (here, closest in attribute space) with unexplored options. However, a prior-free approach need not be pessimistic. We make that point by comparing a max-max searcher to a max-min searcher and find the former behaves similarly to a myopic max-min searcher.

A literature in applied mathematics and machine learning on function optimization posits heuristics which quickly converge to ϵ-optima.[1] Algorithms with good asymptotic performance are useful for many applications but are ad hoc as models of agent behavior. Altering the behavior of such algorithms for a finite number of periods does not affect asymptotic performance. For example, comparative statics or estimates of preferences would differ depending on whether agents are assumed follow gradient descent or simulated annealing.

The present paper is most closely related to literature on Lipschitz function maximization reviewed by Hansen et al. (1992). Several sequential algorithms are designed so as to converge to ϵ-optimal points in no more steps than optimal simultaneous strategies, in the worst case (e.g., see Shubert (1972); Sukharev (1972); Timonov (1977)). Their results correspond most closely to the minimal prior knowledge, zero costs, and dirac-loss function case of our model. We contribute to this literature by allowing for quasiconcave mappings, general function evaluation costs and general loss functions.

Our purpose with these generalizations is to build a model suitable for search theory, which departs from the optimization literature primarily by its focus on optimal search and stopping rules by rational agents. Section 2 and Section 3 describe the model and preliminary concepts. For a general class of history dependent preferences, we characterize all optimal search strategies (Section 4.1, Section 4.2, and Section 4.4). Our main contribution is to use this characterization to derive and show which comparative statics are robust to equilibrium selection (Section 4.5 and Section 4.6). We specialize the model to make sharper comparisons between optimal search dynamics and empirical findings (Section 5). Finally, we extend the model to take a step toward tackling multidimensional search and problems of optimal search space design (Section 6).

[1] See chapter 4.4 of Slivkins (2019) for related work on multi-arm bandits with Lipschitz rewards. See also Aghion et al. (1991) for a perspective on long run optimal learning and experimentation.

2 The Model

Section 2.1 and Section 2.2 describe the model. Section 2.3 discusses interprets some assumptions and simple consequences.

2.1 Preliminaries

There is a set of technologies, represented as a compact set $S \subset \mathbb{R}$. Let Q be the set of potential *quality indices*: each technology $x \in S$ has a quality, $q(x) \in [0,1]$, where $q \in Q$.

There is a searcher. At the start, she knows only the quality of a single technology $x_0 \in S$ and that $q \in Q$. She can learn the quality of other technologies in S through costly *search* (or *experimentation*).

Q is a subset of Lipschitz continuous mappings $S \to [0,1]$ with Lipschitz constant less than or equal to $L > 0$. Continuity of quality indices constitutes the spatial aspect of search over S, as nearby technologies in S cannot vary too drastically in quality.

In each period, $t = 1, 2, 3 \ldots$, the searcher takes one of two kinds of actions. She either experiments with a new technology $x_t \in S$ to learns its quality, $q(x_t)$. Or she concludes her search ($x_t = \emptyset$) and adopts the highest quality technology that she had experimented with so far.

Formally, let $h_t = \{(x_i, z_i)\}_{i=0}^{t-1}$ be a time t partial history when the searcher has not yet concluded its search, with $z_i = q(x_i)$. Let X_{h_t} be the set of technologies discovered at this history. Let $h_t(i)$ denote a time i sub-history for $i \leq t$. Let

$$z^*_{h_t} = \max_{i=0,\ldots,t-1} z_i.$$

If $x_i \in X_{h_t}$ and $z_i = z^*_{h_t}$, then x_i is *an optimal technology at* h_t. Let $X^*_{h_t} \subset X_{h_t}$ denote the set of technologies discovered at this history with quality $z^*_{h_t}$. Let $Q_{h_t} \subset Q$ be the set of quality indices that are *consistent* with what the searcher had observed so far at history h_t. That is, Q_{h_t} is the set of indices q' satisfying $q'(x_i) = z_i$ for all $i = 0, \ldots, t-1$. Let H denote the set of all partial histories h_t where Q_{h_t} is nonempty.

If q is the quality index, the searcher's benefit to adopting technology x is,

$$U(q(x), \max_{y \in S} q(y)),$$

where U is differentiable almost everywhere and continuous. U is non-decreasing in the quality of the adopted technology and non-increasing in the quality of the best technology in S. Moreover, $U_1 \geq -U_2$ almost everywhere, so for any technology choice, translating the quality index q upward always benefits to the searcher.

The searcher's cost of experimenting with technology x in period t is given by

$$C(x, h_t),$$

where $C : S \times H \to \mathbb{R}_{++}$ is bounded away from 0. C may depend on the history of technologies she experimented with so far but not their qualities. For example, the searcher may find it easier to experiment with technologies nearer to those she had explored earlier.

When the quality index is q, the searcher's total payoff after concluding search at history $h_t \in H$ is given by:

$$p(h_t, q) = U(z^*_{h_t}, \max_{x \in S} q(x)) - \sum_{i=1}^{t} C(x_i, h_t(i)).$$

2.2 Searcher's information and objective

The searcher does not know q. She only observes the quality of the technologies that she experiments with. At each history, the searcher worries about the quality index q that is consistent with her earlier experiments but would thereafter minimize her payoff. She seeks a strategy that is *robust* to such outcomes.

We recast the searcher's problem as a dynamic zero-sum game against an imaginary adversary. Each period proceeds in three stages. At history h_t, in the first stage, the searcher acts (by experimenting with a new technology, x_t, or by concluding search, $\emptyset$). In the second stage, the imagined adversary picks a quality index $q' \in Q_{h_t}$ that is consistent with previous experiments at this history. In the third stage, either the quality of this technology is revealed; or the searcher realizes her payoff $p(h_t, q')$, had she concluded search in the first stage.

Let $\sigma : H \to \Delta S \cup \{\emptyset\}$ and $\sigma^A : H \times S \to \Delta Q$ denote strategies of the searcher and its imagined adversary, respectively. The adversary is constrained to choosing strategies which satisfy $\sigma^A(h_t, \cdot) \in \Delta Q_{h_t}$ for every $h_t \in H$; such strategies are said to be *feasible*.

Let $h^{(\sigma,\sigma^A)}$ denote the history after which the searcher concludes search when both agents follow their respective strategies. The payoff to the searcher is $p(h^{(\sigma,\sigma^A)}, \sigma^A(h^{(\sigma,\sigma^A)}))$ and the payoff to the adversary is $-p(h^{(\sigma,\sigma^A)}, \sigma^A(h^{(\sigma,\sigma^A)}))$. If the searcher never concludes search under (σ, σ^A), then let $p(h^{(\sigma,\sigma^A)}, \cdot) = -\infty$.

Given these preferences, if the searcher concludes search at a history h_t, she anticipates the adversary to mix over quality indices in

$$\begin{aligned} &\arg\min_{q' \in Q_{h_t}} \left(U(z^*_{h_t}, \max_{y \in S} q'(y)) \right) \\ =&\arg\max_{q' \in Q_{h_t}} \left(\max_{y \in S} q'(y) \right). \end{aligned}$$

However, if the searcher continues search at h_t, the adversary has incentive to both depress future benefits and prolong costly search along the way.

A strategy $\sigma : H \to \Delta S \cup \{\emptyset\}$ is an *optimal sequential search procedure* for the searcher if (σ, σ^A) is a subgame-perfect equilibrium of the zero-sum game between the searcher and its fictitious adversary for some feasible σ^A.

By restricting attention to subgame-perfect equilibria, we discover search strategies which are robustly optimal at every history. This allows us to study how robustly optimal search unfolds for any $q' \in Q$, and not just for adversarially generated quality indices. Also note, the convention that the searcher knows the quality of only one technology in the beginning is without loss of generality for the analysis of optimal sequential search procedures.[2]

2.3 Discussion

2.3.1 Dependence of payoff on alternatives

One interpretation of the payoff function is that a the searcher is a firm whose future profits depend not only on the quality of its chosen technology but also on the technology chosen by a strong potential competitor. In the worst case, the competitor discovers the best innovation possible, maximizing the gap in the subsequent performance of the two firms.

Alternatively, the searcher could be an online shopper who readily observes some (index of) product characteristics such as prices and appearance. By clicking on the product and reading its description and reviews, she learns her value for the purchase. But she has reference dependent preferences and experiences disutility from later learning that there was a better deal that she missed out on.

2.3.2 An independent outside option

The model accommodates the interpretation that the searcher chooses between the best explored technology *or* an outside option that gives a payoff of u_0 independent of the realized quality index. For example, U may describe her anticipated payoff if she launches a startup in a space where one of the competitors may discover a better technology. Alternatively, the searcher could avoid creating a startup altogether and keep her current job for a guaranteed payoff of u_0.

To see why, note that such a searcher's benefit to concluding search at history h_t would be

$$U'(z^*_{h_t}, \max_{y \in S} q(y)) \equiv \begin{cases} u_0, \text{ if } \min_{q' \in Q_{h_t}} U(z^*_{h_t}, \max_{y \in S} q'(y)) < u_0. \\ U(z^*_{h_t}, \max_{y \in S} q(y)), \text{ otherwise.} \end{cases}$$

[2]The results are unchanged even if we assume the searcher does not know the payoff to any technology to begin with. We sometimes adopt this convention when convenient.

Since U is nondecreasing in the first argument and non-increasing in the second argument, the same is true for U'; and $U'_2 \le U'_1$, as well. Therefore, the original model with U' in place of U is equivalent to the model with an outside option.

3 Simultaneous Search Procedures

Here we introduce terminology for certain subclasses of the searcher's strategy space that will be useful for the analysis in Section 4.

A simultaneous search procedure is one that depends only on the technologies that the searcher explored so far but not their qualities. Intuitively, for a searcher who uses a simultaneous search procedure, it is immaterial whether she learns the results of her experiments as she goes or only at the end of her search.

Formally, two partial histories $h'_t, h''_t \in H$ *differ only by quality* if the same sequence of technologies are explored in both histories, i.e., $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1}$ and $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1}$. A strategy σ is a *simultaneous search procedure* if $\sigma(h'_t) = \sigma(h''_t)$ for any $h'_t, h''_t \in H$ that differ only by quality.[3]

We similarly define the notion of a strategy which becomes a simultaneous procedure upon reaching a certain history. A strategy σ is a *simultaneous search procedure at* h_t if $\sigma(h'_t) = \sigma(h''_t)$ whenever $h'_t, h''_t \in H$ differ only by quality and h_t is a sub-history of both h'_t and h''_t.

Let Γ_{h_t} denote all simultaneous search procedures at h_t. Consider a modification of the sub-game starting at history h_t where the searcher is constrained to the strategy space Γ_{h_t}; this is the *simultaneous search game at* h_t. A strategy σ is an *optimal simultaneous search procedure at* h_t if $\sigma \in \Gamma_{h_t}$ and if there exists some strategy σ^A for the adversary such that (σ, σ^A) is a Nash equilibrium of the simultaneous search game at h_t.

Finally, we can define a strategy that solves for an optimal simultaneous search procedure at every history and employs that strategy for a single period. Let σ^*_{s,h_t} denote some optimal simultaneous search procedure at h_t. A strategy σ *follows an optimal simultaneous search procedure at every history* if $\sigma(h_t) = \sigma^*_{s,h_t}(h_t)$ for all $h_t \in H$. Note that such a strategy is typically not itself a simultaneous search procedure at any history.

3.1 Following optimal simultaneous search procedures can be sub-optimal

By following an optimal simultaneous search procedure at every history, the searcher may disregard opportunities to acquire valuable information and end up searching more than is

[3] Simultaneous search procedures may depend on the order in which a set of technologies is searched.

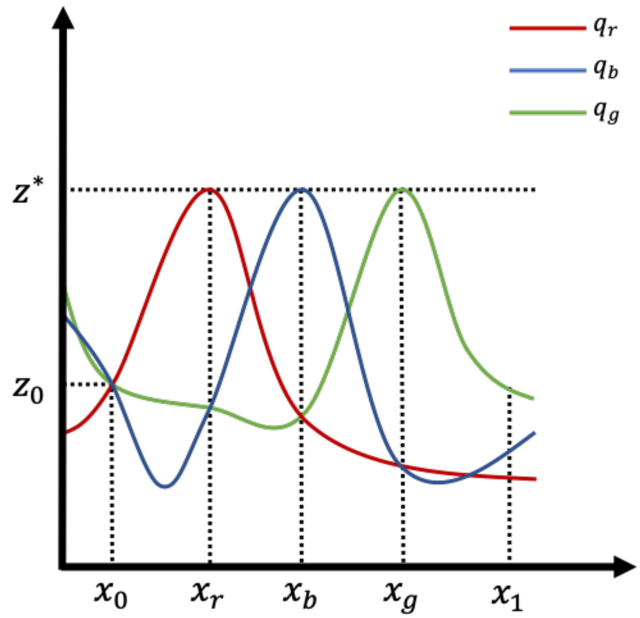

Figure 1: The quality indices q_r, q_b, and q_g attain their maxima at x_r, x_b, and x_g respectively.

necessary.

For example, let $Q = \{q_r, q_b, q_g\}$ as pictured in Figure 1. Suppose that search costs are constant: $C(x, h_t) = \epsilon$ for some $\epsilon > 0$, for all $x \in [0, 1]$ and $h_t \in H$. Suppose U is strictly increasing in the quality of the adopted technology. The searcher starts by observing the quality of technology x_0. Since $q_r(x_0) = q_b(x_0) = q_g(x_0) = z_0$, the searcher cannot narrow down the set of possible quality indices at history $\{(x_0, z_0)\}$.

For ϵ sufficiently small, the optimal simultaneous procedure is to experiment with x_r, x_b and x_g in any order. One of these technologies is guaranteed to be of quality z^*. Experimenting with any other technologies is clearly wasteful. And leaving one or more of x_r, x_b or x_g unexplored lowers the searcher's worst-case payoff. For example, she is worse off if she only searches x_b and x_g, and q_r turns out to be the true quality index.

On the other hand, no optimal sequential search procedure starts with x_r, x_b or x_g. For instance, if the searcher starts with x_r, the true quality index is not q_r in the worst case, but either q_b or q_g. Since $q_b(x_r) = q_g(x_r)$, the searcher entertains both possibilities. It would take her two more tries to guarantee finding the maximum of the true quality index. However, if the searcher starts by exploring x_1, she immediate identifies the true quality index, since $q_r(x_1) \neq q_b(x_1) \neq q_g(x_1)$. In the next step, she would choose the technology that maximizes that index. This strategy guarantees z^* with two costly searches, rather than three.

4 Optimal Sequential Search Procedures

4.1 Optimal sequential search with minimal prior knowledge

Suppose that a searcher knows only that the quality index cannot vary too drastically for nearby technologies; she knows nothing else about the shape of this index.

Formally, let Q^{MP} denote the set of all Lipschitz continuous mappings $S \to [0, 1]$ with Lipschitz constant less than $L > 0$. The searcher has *minimal prior knowledge* if $Q = Q^{MP}$.

The first result establishes that, unlike the example in Section 3.1, following optimal simultaneous search procedures is optimal when the searcher has minimal prior knowledge.

THEOREM **1.** *Suppose $Q = Q^{MP}$. Then a search strategy σ is an optimal sequential search procedure if and only if it follows an optimal simultaneous search procedure at every history.*

The proof of Theorem 1 is in Appendix B.1, while an intuition is given in Section 4.4.

This result shows that optimal sequential search reduces down to a type of greedy simultaneous search. The searcher solves for an optimal simultaneous search procedure at every period, but only executes the fist step of that plan. If the realized quality from that experiment is what she expected, she continues to carry out the previously calculated optimal simultaneous search procedure. If quality ends up being lower or higher than she anticipated, (i.e., the worst-case did not actually materialize), she formulates a new plan.

4.2 Optimal sequential search over quasiconcave quality indices

Let Q^{QC} denote the set of all quasiconcave Lipschitz continuous mappings $S \to [0, 1]$ with Lipschitz constant less than $L > 0$. The searcher *additionally knows that quality is quasiconcave* if $Q = Q^{QC}$. Compared to the minimal prior knowledge case, such a searcher also knows that there is an ideal technology and that technologies that are farther from this ideal point are of weakly lower quality.

The next result gives an analogue to Theorem 1 when the searcher additionally knows that quality is quasiconcave.

THEOREM **2.** *Suppose $Q = Q^{QC}$. Then a search strategy σ is an optimal sequential search procedure if and only if it follows an optimal simultaneous search procedure at every history.*

4.3 Additional Bounds on Q

If the searcher knows that she prefers higher values (e.g., picture quality) or lower values (e.g., price) of the dimension along which she searches, it may make sense to impose certain bounds on Q. For example, suppose someone searching a lineup of TVs by price perceives the 'quality' of the product to be an increasing function of the initially unobserved picture quality minus the observable price. Such a searcher may believe that TVs above (or below) a certain price would not be great bargains. Q may then be the set of all functions in Q^{MP} that are lie below $f(x) = 1 - (0.5 - x)^2$.

Theorem 1, Theorem 2, and the remaining results of the paper go through even when such bounds are placed on Q. Of course, such bounds will affect optimal search procedures. For example, all else equal, the searcher may avoid searching areas of S where the maximal and minimal values of q are known ex-ante to be close. Searching elsewhere would be more informative.

4.4 Optimal simultaneous search

Theorem 1 and Theorem 2 reduce the problem of sequential search to solving for an optimal simultaneous search procedure at every history. Here we characterize optimal simultaneous search procedures.

Let $q^u_{h_t}$ denote the upper envelope of quality indices in Q_{h_t} for any $h_t \in H$, and let $q^A_{h_t} = \min\{q^u_{h_t}, z^*_{h_t}\}$. Let $\sigma^A_d(h_t, x) = q^A_{h_t}$ for all $x \neq \emptyset$, and $\sigma^A_d(h_t, \emptyset) = q^u_{h_t}$.

LEMMA **1.** *Suppose that $Q = Q^{MP}$ or $Q = Q^{QC}$, and let $\sigma \in \Gamma_{h_t}$. Then σ is an optimal simultaneous search procedure at h_t if and only if σ is a best-response to σ^A_d in the simultaneous search game at $h_t \in H$.*

When it comes to keeping down the benefits that the searcher ultimately receives upon concluding search, the adversary's strategy serves two purposes. First, the adversary does not let the quality of any searched item exceed $z^*_{h_t}$ because that improves the quality of the searcher's chosen item more than it increases the outcome of the potential best alternative; the trade-off would be in the searcher's favor. Next, the adversary otherwise picks the quality to be as high as possible, since this would improve the size of the best potential alternative without increasing the quality of the searcher's chosen outcome.

The same intuition for keeping down benefits extends from the simultaneous search game to the sequential search game. But a new incentive for the adversary in the sequential search game is to keep the searcher engaged in search by giving little information on where good outcomes lie. Precisely the same strategy used to keep down searcher benefits is useful for this purpose as well. Seeing higher or lower qualities than $z^*_{h_t}$ would allow the searcher to focus experimentation in more promising regions of the search space. A strategy that shows neither good nor bad news keeps would induce the searcher to continue to search broadly, on a wild-goose chase.

This link between the dual objectives of keeping down benefits and suppressing information is reflected in the link between optimal simultaneous and sequential search established

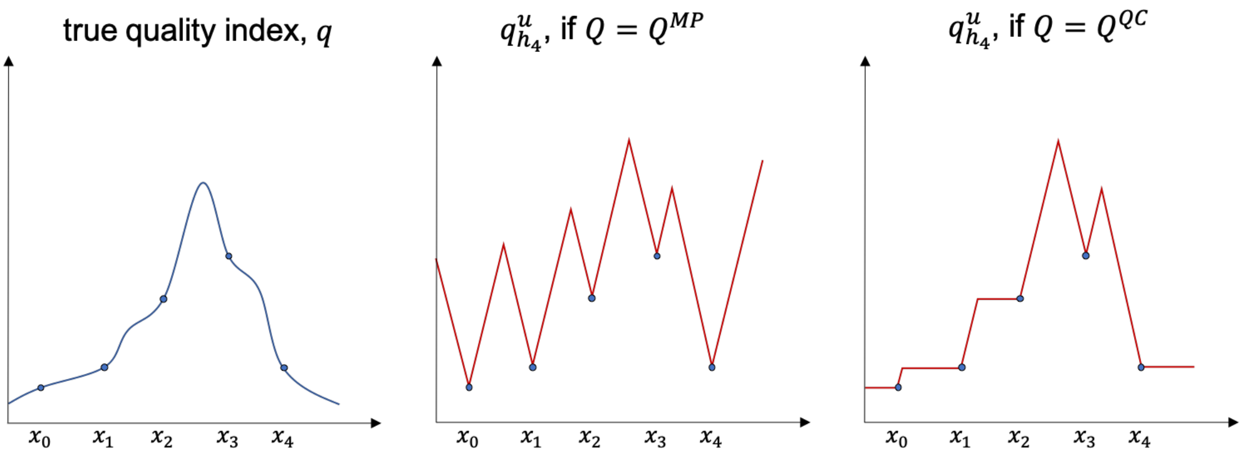

Figure 2: The figure on the left displays the true quality index. At history h_4, the searcher has observed the qualities at x_0, x_1, x_2, x_3 and x_4 in that order. The figure in the middle plots $q^u_{h_4}$, the upper envelope of feasible quality indices at history h_5, when $Q = Q^{MP}$. The figure on the right plots $q^u_{h_4}$ when $Q = Q^{QC}$. Note that when $Q = Q^{MP}$, the searcher may be interested in experimenting with technologies to the right of x_4 to rule out the possibility of finding something of high quality there. However, in the case where $Q = Q^{QC}$, the searcher is guaranteed not to find any breakthrough technologies in this region.

in Theorem 1 and Theorem 2.

Notwithstanding this similarity, the dynamics of optimal sequential search play out very differently when $Q = Q^{MP}$ versus when $Q = Q^{QC}$. That is because at any history, the worst-case outcome from the searcher's perspective can look quite different depending on whether or not she believes the quality index to be quasiconcave. Different perceptions of the worst-case would then feed into different choices for where to search next. Figure 2 gives an example of how the searcher's perceived worst-case can differ depending on Q, starting from the same history.

4.5 Complexity and propensity to search

A basic question is when the searcher explores any technologies at all and when she continues with the status quo. For example, understanding in which settings there may or may not be endogenous information acquisition is a first step to thinking about when such efforts should be subsidized. Here we explore how the propensity to search varies with search complexity, measured by L.

Search complexity can be a literal description of how hard it is to discover a relatively good outcome. For example, a pharmaceutical company may experiment with a compound whose efficacy and safety is typically very sensitive to dosage. In the worst case, finding the right dosage may be like finding "a needle in a haystack". Similarly, a software firm may be competing to make a product where slight differences in design significantly affect user experience and market share.

The first result is that search duration is non-monotonic in the search complexity. In particular, the searcher does not explore at all if search is sufficiently complex or simple.

PROPOSITION **1.** *Suppose $Q = Q^{MP}$ or $Q = Q^{QC}$, and $|S| = \infty$. Let $h \in H$. There exist $\underline{L}, \overline{L} > 0$ such that if search complexity $L < \underline{L}$ or $L > \overline{L}$, and if Q_h is nonempty, the optimal sequential search procedure at h concludes search immediately.*

The intuition for this result is that when L is sufficiently small, the value of search falls since the searcher expects the quality index to be relatively flat and close to x_0.

If, on the other hand, L is sufficiently large, then the searcher is discouraged for a different reason: she may spend a lot of resources searching and still come nowhere near the peak outcome. This intuition hinges on the quality index being bounded, i.e., $q \in [0, 1]$. When there are many $x \in S$ for which the quality could potentially attain the upper bound, then one search gives little or no additional information about an unexplored technology.

If, instead, the highest achievable quality falls short of the bound, searches are much more informative about unexplored technologies. Suppose, in addition, that there are *decreasing returns to reducing the peak of the quality index*, i.e., $U_{22} < 0$. For example, a shopper is more concerned with missing out on larger deals than smaller ones. Then the propensity to search is non-decreasing in search complexity.

PROPOSITION **2.** *Suppose $Q = Q^{MP}$ or $Q = Q^{QC}$, and there are decreasing returns to reducing the peak of the quality index. Let $0 < L' < L''$, $h_0 = \{(x_0, z_0)\} \in H$, and suppose that $q_h^u < 1$ when search complexity is L''. If search concludes immediately under some optimal sequential search procedure when search complexity is L'', then search concludes immediately under any optimal sequential search procedure when search complexity is L'.*

The proof follows a dominance argument. Suppose a particular sequence of searches brings the searcher closer to the peak when the complexity is low, and searching is worth the cost. The same sequence of searches will bridge a larger gap between the peak and the quality of the best technology discovered so far when the search complexity is higher.

Taken together, Proposition 1 and Proposition 2 suggest that the propensity to search takes an inverse-U shape in complexity when there are decreasing returns to being near the peak of the quality index. In other words, inducing exploration requires both the promise of being able to approach the peak of the quality index with some effort and the threat of being currently far from it.

4.6 Myopic search strategies

Here we characterize myopic search procedure.

This is helpful in understanding how search in a social learning setting compares to search when information externalities are internalized. Myopic search in this case can be seen as the outcome of decisions by a sequence of short-lived agents who learn from the experiments of their predecessors.

Moreover, structurally estimated models of spatial search assume, out of necessity, that searchers look ahead only one period (e.g., Hodgson and Lewis (2020)). The differences between myopic and optimal search in the present model may be suggestive of how assuming bounded rationality biases estimation results.

A myopic strategy is one where which takes the best step (e.g., either concluding search or searching some new technology) at each history if the searcher had to conclude search by the subsequent period.

Formally, let $M > 0$ be large enough such that if the cost of search for any $x \in S$ is at least M, an optimal search procedure would conclude search immediately at any $h_t \in H$. For any $h_t \in H$, let $\sigma_{h_t}^{SL}$ denote a optimal sequential search procedure if costs where given by C^{SL}, where $C^{SL}(x, h_t) = C(x, h_t)$ for any new technology $x \in S$, and $C^{SL}(\cdot, h_{t+i}) = M$ for all $i > 1$. A *myopic strategy* σ^M is a strategy that follows $\sigma_{h_t}^{SL}$ at every history $h_t \in H$.

4.6.1 Propensity to Search

A simple observation is that myopic searchers have a lower propensity to continue search than forward looking searchers. More formally, at any $h_t \in H$, if there exists an optimal sequential search procedure σ such that $\sigma(h_t) = \emptyset$, then there is a myopic strategy σ^M such that $\sigma^M(h_t) = \emptyset$. This follows by the optimality of σ and the definition of σ^M: since $\sigma^M \neq$, the payoff to continuing search at h_t for one period and concluding search subsequently is weakly greater than concluding immediately.

4.6.2 Search Location

To characterize where myopic searchers would look, we introduce some terminology.

For any $h_t \in H$, the *search window*, $W_{h_t} \subset S$, is the set of technologies x for which $q_{h_t}^A(x) = z_{h_t}^*$. Outside of the search window, the largest realizable quality is strictly below $z_{h_t}^*$ for any $q \in Q_{h_t}$.

Let $x = \arg\max_{y \in S} q_{h_t}^u(y)$, supposing that the maximizer is unique. Suppose moreover that $x \notin X_{h_t}$, so that there is still a positive benefit to searching. Let $x_r = \min\{y \in X_{h_t} | y > x\} \cup \{\max S\}$, i.e., x_r is the closest discovered technology to the right of x and $\max S$ if

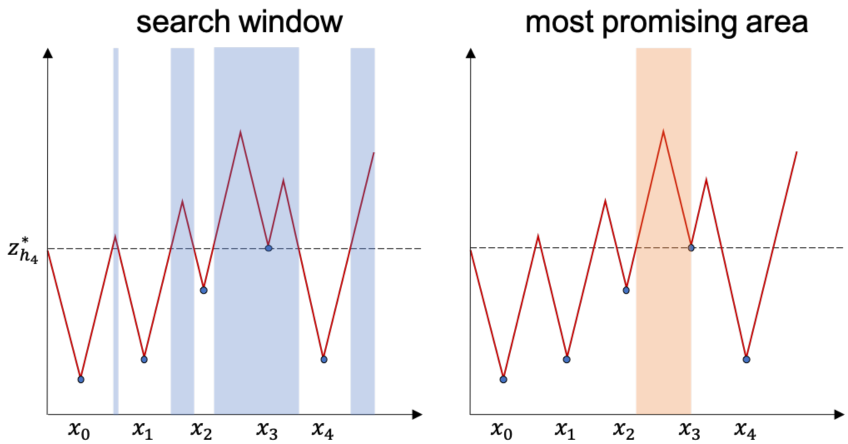

Figure 3: Continuing the example from Figure 2, both figures plot $q^u_{h_4}$ when $Q = Q^{MP}$. The figure on the left highlights the search window whereas the figure on the right highlights the most promising area.

there is no such technology. Similarly, let $x_l = \max\{y \in X_{h_t} | y < x\} \cup \{\min S\}$. Define $[x_l, x_r] \cap W_{h_t}$ to be the *most promising area* of S at h_t (see Figure 3).

The next proposition gives a general characterization of where myopic search occurs.

Proposition **3.** *Let $Q = Q^{MP}$ or $Q = Q^{QC}$, and let $h_t \in H$. For any myopic strategy, σ^M, either $\sigma^M(h_t) = \emptyset$, or $\sigma^M(h_t)$ is in most promising area of S at h_t.*

Intuitively, any future search in the most promising area results in one of two possible outcomes: either the searcher learns that the highest possible quality is smaller than expected or she finds a better technology than previously discovered. Outside of the most promising area, finding technologies that are worse than what was previously discovered does not change the searcher's perception of what the highest possible quality could be. Searches in this region are of no value to a myopic searcher who does not expect to find something of higher quality than $z^*_{h_t}$.

Unlike myopic strategies, optimal sequential search procedures do not always search in the most promising area of S at every history.

For example, suppose the cost of experimenting with a set of promising technologies is prohibitively high for a new startup with limited resources. However, experimenting first with a more basic but less promising technology makes exploring the more promising technologies accessible (e.g., by attracting attention from investors or by training the designers). A forward-looking startup may start with the basic technology to unlock lower costs of future experimentation. A myopic startup that wishes to sell for the highest valuation after one period may find it worthwhile to experiment with the most promising technologies directly.

This is an example where costs depend on the history of searches. Section 5 considers the case of history independent costs where differences between myopic and optimal strategies persist.

5 Fixed Costs Model

The model of spatial search is quite general and can be adapted to a variety of functional forms. Here, we restrict attention to the natural classes of history-independent cost functions to derive a richer set of comparative statics and explicit solutions to the searcher's problem.

Let $f : S \to \mathbb{R}_+$ be a continuous function, bounded below by $c > 0$. A *fixed cost model* is one in which $C(x, \{(x_i, z_i)\}_{i=0}^t) = f(x)$. A special case is the *constant fixed cost model*, where f is constant.

5.1 Impact of news on search dynamics

Fix the true quality index q and a search strategy. Consider a partial history $h_{t+1} \in H$ where the searcher does not conclude search, and let h_t be its time t sub-history. The searcher learns *good news* at h_{t+1} if $z_{t+1} > z^*_{h_t}$. Similarly, the searcher learns *bad news* at h_{t+1} if $z_{t+1} < z^*_{h_t}$.

5.1.1 Search step size

The next proposition describes the impact of sufficiently bad news on step size in an optimal search procedure. It captures the intuitive idea that bad news should drive the searcher to avoid the surrounding area and search elsewhere.

PROPOSITION **4.** *Suppose $Q = Q^{MP}$ or $Q = Q^{QC}$, and consider a fixed costs model. Let σ be an optimal sequential search procedure and consider any history $h_t \equiv \{(x_i, z_i)\}_{i=0}^t \in H$ at which $\sigma(h_t) \neq \emptyset$ and $\min_{q \in Q_{h_t}} q(\sigma(h_t)) > 0$. Let x be the closest technology to $\sigma(h_t)$ among those that had been discovered so far. Then there exists $z' < z^*_{h_t}$ such that:*

1. *$Q_{h'_{t+1}}$ is non-empty, where $h'_{t+1} \equiv \{(x_0, z_0), \ldots, (x_t, z_t), (\sigma(h_t), z')\}$.*
2. *If $z_{t+1} \leq z'$, then either $\sigma(h_{t+1}) = \emptyset$ or $|\sigma(h_{t+1}) - \sigma(h_t)| > |\sigma(h_t) - x|$.*

The second part of Proposition 4 says that after sufficiently bad news, an optimal search procedure either concludes or jumps beyond the nearest technology to another part of the search space. The first part of Proposition 4 says that, at every history, it is possible to hear such sufficiently bad news.

Notice that with bad news, the best technology discovered so far remains unchanged, but not so with good news. Depending on the utility function and the history of technologies searched, the searcher's reaction to good news varies.

The searcher may care less about being near the top as the quality of her best discovery rises. In that case, she may search less carefully in reaction to good news and make larger jumps in search space.

Alternatively, the searcher may care more about being near the top as the quality of her best discovery rises. This happens when good news signals higher stakes. For example, if a firm's search for an efficient flying car technology produces poor outcomes, the firm might come to believe the market for such cars will never be large. But if the firm chances upon a good design, it may come to believe that there is potentially a large market and the winner in the race to build the best flying car would capture a disproportionate share. Such a searcher may carefully optimize her design by searching the space near the best discovery so far before directing her attention elsewhere.

It is simple to give examples where these intuitions hold, but difficult to generalize them. Whether or not the searcher reacts to good news in a certain way also depends finely on her history of discoveries interacts and cost function.

5.1.2 Triangulation

The impact of good or bad news on search dynamics is more striking when the quality index is known to be quasiconvex. The searcher can use *any* good or bad news to significantly narrow down the space over which future searches will occur.

Starting at some history $h_t \in H$, a searcher *never searches in some set* $T \subset S$ *under strategy* σ if for any $h \in H$ such that $h(t) = h_t$, either $\sigma(h) \in S/T$ or $\sigma(h) = \emptyset$.

PROPOSITION **5.** *Let* $Q = Q^{QC}$ *and consider a fixed costs model. Suppose the searcher uses some optimal sequential search procedure* σ*, and let* $h_t = \{(x_i, z_i)\}_{i=0}^{t} \in H$ *be a history at which* $\sigma(h_t) \neq \emptyset$*.*

1. *Suppose the searcher learns good news at* $\sigma(h_t)$*, i.e.,* $z_{t+1} > z^*_{h_t}$*.*
   (a) *If* $\sigma(h_t) > x_t$*, the searcher will never search left of* $x_t + \frac{1}{L}(z_{t+1} - z^*_{h_t})$*.*
   (b) *If* $\sigma(h_t) < x_t$*, the searcher will never search right of* $x_t - \frac{1}{L}(z_{t+1} - z^*_{h_t})$*.*
2. *Suppose the searcher learns bad news at* $\sigma(h_t)$*, i.e.,* $z_{t+1} < z^*_{h_t}$*.*
   (a) *If* $\sigma(h_t) > x_t$*, the searcher will never search right of* $\sigma(h_t) - \frac{1}{L}(z^*_{h_t} - z_{t+1})$*.*
   (b) *If* $\sigma(h_t) < x_t$*, the searcher will never search left of* $\sigma(h_t) + \frac{1}{L}(z^*_{h_t} - z_{t+1})$*.*

Good and bad news events create search dynamics that resemble a 'funnel shape' or 'triangulation' or 'zooming in on search': the walls close in on the space of products or technologies over which search unfolds. This dynamic of search broadly covering attribute space at first and then narrowing in on a particular region was observed in both Blake et al. (2016) and Bronnenberg et al. (2016). Learning and the belief in a 'sweet spot' are sufficient, and natural, conditions for generating such dynamics.

5.2 An example with constant fixed costs

Here, we describe an algorithm for computing optimal sequential search procedures for a constant fixed cost model and a connected search space. We then compare optimal search to myopic search in this special case.

Let $S = [0, 1]$ and $C(\cdot, \cdot) = c > 0$. Suppose $Q \in \{Q^{MP}, Q^{QC}\}$.

5.2.1 Optimal search in the constant fixed cost model

We describe the algorithm in words and leave the formal description for Appendix A.

Recall the definition of a search window in Section 4.6.2. At any history, $h_t \in H$, the searcher considers the best set of technologies to search, anticipating the adversary to best respond. When costs are history-independent, there is no value in searching outside the search window. Such a search would return a lower payoff than $z^*_{h_t}$ without affecting the costs of future searches.

However, searching inside the search window produces two possible outcomes: either the searcher discovers a better quality technology or the technology quality is low which constrains how good nearby unexplored technologies can be. In the worst case, the latter happens, bringing the searcher's technology closer to the best possible technology.

Indeed, by Lemma 1, the searcher expects to discover a technology of quality $z^*_{h_t}$ for every search. This implies that the searcher expects the search window to remain unchanged for the duration of search.

Note that if $Q = Q^{QC}$, by Proposition 5, the search window at any history, $h_t \in H$, is a single interval. If $Q = Q^{MP}$, the search window is a disjoint union of intervals. Let Y_{h_t} be the set of endpoints of the disjoint intervals that comprise the search window and technologies in X_{h_t} which fall in the search window.

At the midpoint of two adjacent, path-connected technologies $y, y' \in Y_{h_t}$, $q^u_{h_t}$ has a local maximum (see Figure 3). Clearly, the size of that local maximum is increasing in the distance between y and y'. Searching a technology $y'' \in (y, y')$ returns a quality of $z^*_{h_t}$ in the worst-case. The searcher therefore anticipates $q^u_{h_{t+1}}$ have peaks in (y, y'') and (y'', y'), both smaller than the peak $q^u_{h_t}$ has in (y, y').

The searcher's goal is to reduce the height of the tallest peak, upon concluding search: this would be the quality of the best unexplored alternative in the worst case. If the searcher were to conclude search after k more periods, a simple geometric intuition suggests that, at h_t, she should plan to pick her k search points so that $Y_{h_{t+1}}$ partitions $W_{h_{t+k}}$ as evenly possible. This way, the largest width between adjacent points in $Y_{h_{t+k}}$ is minimized, making the global maximum of $q^u_{h_{t+k}}$ as small as possible on the equilibrium path.

This logic pins down the location of searches conditional on k. The benefits function U plays a role only to determine how large she should plan for k to be at h_t, i.e., how many additional searches are worthwhile. Clearly, searching too much would eventually push payoffs below the payoff to concluding search immediately. This serves as an upper bound on k. The searcher can compute the worst-case payoff to planning for fewer searches for all smaller values of k and proceed with the plan that returns the highest payoff.

Once the optimal k is determined and the k technologies that most evenly partition $W_{h_{t+k}}$ are determined, any search that picks one of these k technologies is part of an optimal sequential search procedure.

5.2.2 Myopic search in the constant fixed costs model

In the case of constant fixed costs, one natural myopic strategy is simple to describe. Let $h_t \in H$. Let $x \in \arg\max q^u_{h_t}$, and let h_{t+1} be the history after h_t where x is searched and found to have quality $z^*_{h_t}$.[4] Define

$$\sigma^M(h_t) \equiv \begin{cases} \emptyset, \text{ if } p(h_t, q^u_{h_t}) > p(h_{t+1}, q^u_{h_{t+1}}) \\ x, \text{ otherwise} \end{cases}$$

It is clear that σ^M is a myopic strategy. In the event search does not conclude immediately, σ^M explores the technology that leads to the one-period greatest decrease in the quality of the best available alternative, on the equilibrium path (Lemma 1).

The technology explored by σ^M is farther away from Y_{h_t}, in terms of the usual set distance, than any other technology in the search window. In this sense, this myopic strategy ventures into a more unexplored part of the search space than any optimal sequential strategy.

Suppose for a moment that a searcher following an optimal sequential search procedure receives bad news and jumps farther away in her next search than previously anticipated. An outside observer who knows that the searcher saw bad news but mistakenly perceives the searcher to follow the described myopic strategy may underestimate how bad the news was. Larger steps away from Y_{h_t} can be rationalized for a myopic searcher without significant bad news.

[4] It is easy to verify that $Q_{h_{t+1}} \neq \emptyset$

5.2.3 Optimistic search in the constant fixed cost model

As an aside, a critique of worst-case analysis is that real-world decision-makers may not be so pessimistic. This raises the question of how to capture the behavior of more optimistic decision-makers. In the case of fixed constant costs, it is simple to fully characterize the behavior of the most optimistic agent, i.e., one who believes, at every history $h_t \in H$, the best outcome in Q_{h_t} will obtain given her actions.

At any $h_t \in H$ at which such an agent does not conclude search, she will next search an $x \in \arg\max q^u_{h_t}$, just as in σ^M. The only difference between the optimistic searcher and the myopic (pessimistic) searcher is in their stopping rules. The former expects this search to return quality $q^u_{h_t}(x)$ and for search to be unnecessary in all future periods. The latter expects this search to return quality $z^*_{h_t}$ and by definition will not search further.

All optimal search procedures for the optimistic searcher take this form. Choosing an $x \notin \arg\max q^u_{h_t}$ is sub-optimal for the optimist. Therefore, optimistic searchers always venture into more unexplored areas of the search space than pessimistic searchers.

Outside of constant fixed costs, location choices for myopic and optimistic search need not coincide. Indeed, even in a more general fixed costs model, optimistic searchers may search outside of the most promising area entirely, if search in this region is significantly more expensive than search in other regions. Recall, however, that search outside of the most promising region is unhelpful to a pessimistic, myopic searcher (Section 4.6.2).

5.3 Selection rules for the fixed cost model

In general, optimal sequential search procedures are not unique in a fixed cost model, as can be seen from the last step in Section 5.2.1. It is of interest (e.g., for inferring preference parameters from search data) to consider various natural selection rules to narrow down the set of optimal search procedures.

5.3.1 Travel costs model

Let $g : \mathbb{R}_+ \to \mathbb{R}_+$ be a strictly increasing, strictly convex function with $g(0) = 0$. A *travel costs model* is one in which $C(x, \{(x_i, z_i)\}_{i=0}^t) = f(x) + g(|x - x_t|)$.

In the travel costs model, the cost of search includes a fixed component, which depends on the technology or product being searched, and a travel component, which depends on the distance between the previous and current searched technologies. This can describe, for example, the form of search costs of an online shopper who is navigating a price-ordered list within some product category. Searching a given item involves the fixed cost of clicking and waiting for the page to load. Jumping to a farther away item in price-space involves more scrolling or navigating to a new page.

Introducing travel costs drastically reduces the multiplicity of optimal search procedures, since maintaining the direction of search and exploring nearer experiments first saves on costs.

We can use this idea while maintaining the tractability of the fixed cost model by considering the case of vanishing travel costs.

Let $c, \{\alpha_n\}_{n=1}^{\infty} > 0$, and let $C_n(x, h_t) = c + \alpha_n g(n|x - x_t|)$, where $\alpha_n g(n|x - x_t|) \to 0$ as $n \to \infty$. In effect, travel costs go to zero while becoming steeper. For large n, optimal sequential search procedures converge to the set of the optimal sequential search procedures in a fixed costs model.

Let T_{h_t} be the set of technologies that the searcher would explore after h_t along some equilibrium path in some optimal sequential procedure in a fixed cost model. Optimal sequential search procedures converge to procedures where the searcher starts on a the *hamiltonian cycle* through T_{h_t} that minimizes the maximum step size between all future experiments along the anticipated equilibrium path.[5] Because hamiltonian cycles minimizing the max step size are generically (in T_{h_t}) unique, this selection criterion gives a sharp prediction for search dynamics in the fixed cost model.

This selection rule captures settings where travel costs are a secondary concern to the costs of conducting experiments themselves.

5.3.2 Search distractions

One can also consider a modification of the fixed cost model where search becomes prohibitively expensive with some small probability at every step. This can be interpreted as a distraction leading to a premature conclusion of search. As this probability vanishes, an optimal sequential search procedure with search distractions converges to an optimal sequential search procedure where at each history h_t, the searcher chooses a technology that leads to the biggest drop in $\max_{x \in s} q^u_{h_t}(x)$.

6 Extensions to Multidimensional Search

Often, technologies or products have multiple observable attributes. For example, online marketplaces may make the resolution, size and price of TVs readily observable before shoppers have to click on the image of the product to learn more. In effect, the search space is multidimensional.

[5]Given a set of n points $T \subset R$ and an initial point, $t_0 \in T$, a *hamiltonian cycle* is a sequence $(t_0, t_1, \ldots, t_n)$ that includes each other point in T exactly once. Since travel costs are convex and grow steeper in n, the largest step size $|t_{i+1} - t_i|$ dominates the aggregate travel costs.

In some settings, searchers may understand how attributes contribute to the quality of an item. In other settings, they might only be able to learn the quality of an item as a whole.

Take a wine drinker who tastes different wines (at a cost) to discover one that she likes. If she is a novice, she may be able to tell only how much she likes or dislikes what she is tasting. If she enjoys a particular wine, she may try others that appear similar. An expert, on the other hand, may understand how she favors each dimension. She may ask to try something drier but with the same acidity.

This suggests two natural ways of extending the model to multidimensional search. Section 6.1 considers a model where the searcher learns her preferences for attributes. Section 6.2 considers a model where the searcher only learns her value of an item as a whole, though the item has multiple observable attributes. Both sections partially extend the characterizations of optimal sequential search.

These extensions may be helpful in applying the model to multidimensional search data. To that end, a desideratum would be that optimal search is easy to simulate. Section 6.3 describes a polynomial-time algorithm in the case of constant fixed search costs.

Online marketplaces can also make use of such an algorithm when making decisions on how to order items in search space. Section 6.4 discusses this application.

6.1 Learning about Attributes

We define a *attribute learning model*, which is one multidimensional extension of the baseline model.

Suppose items have k attributes, e.g., the size, resolution and brand of TVs. Let S_i be the one-dimensional sets of values that the ith attribute can take. The search space is $S \equiv S_1 \times \ldots \times S_k$.

A attribute quality index is a mapping $\kappa : S \to [0,1]^k$, where $\kappa \in (Q^{MP})^l \times (Q^{QC})^{k-l}$ for some $l \in \{0, 1, \ldots, k\}$.[6] Each time the searcher explores an item, she learns her value for each dimension of the object, e.g., how much she likes a large versus small TV, how much she values more resolution, etc.

The quality index aggregates the qualities of the separate dimensions: for any $x \in S$, $q(x) = f \circ \kappa(x)$, where $f : [0,1]^k \to [0,1]$ is increasing in all its arguments. We say *attributes are substitutes* if f is submodular.

S is a *rectangular search space* if for every $x, y \in S$, there exists some $z \in S$ such that $\kappa(z) = \kappa(x) \vee \kappa(y)$. In words, if the searcher likes some attributes of one item and some of another, she can always find an item in search space that has all the better of every attribute.

To simplify the analysis, we assume that the searcher can conclude search by taking a

[6] Results carry over when the Lipschitz constants are differ across attributes.

previously unexplored option without paying an additional search cost. For example, if she considers a 65 inch 4k TV and a 75 inch 5k TV and learns that she likes the smaller size but a higher resolution, she can purchase a 65 inch 5k TV without incurring additional search costs.[7]

The attribute learning model is otherwise identical to the baseline model. We have the following corollary to Theorem 1 and Theorem 2.

COROLLARY **1.** *Consider an attribute learning model where S is a rectangular search space and attributes are substitutes. Then a search strategy σ is an optimal sequential search procedure if and only if it follows an optimal simultaneous search procedure at every history.*

Corollary 1 implies that previous comparative statics results and algorithms can be readily adapted to the attribute learning model. For example, triangulation in search space occurs along those dimensions of κ which the searcher perceives as being quasiconcave.

When $S = [0, 1]^k$, the same algorithm described in Section 5.2.1 can be applied to choose search location for each attribute. A similar procedure works when S is finite as well. Section 6.3 describes an algorithm for an alternative generalization to multidimensional search and discusses important properties; those results readily apply to the attribute learning model as well.

The assumptions that the search space is rectangular or that attributes are substitutes rule out some interesting applications. For example, it may be possible to purchase a powerful but bulky computer or a slower but more portable model. However, there may not exist a very fast and perfectly portable computer on the market. Similarly, a searcher may like a bubblegum pink convertible but may find the same color to be distasteful for a station wagon. Our results do not cover such cases, though it would be valuable to explore results in this direction.

6.2 Optimal Multidimensional Sequential Search

Theorem 1 generalizes to the multidimensional case:

THEOREM **1′.** *Suppose $Q = Q^{MP}$ and S is a compact subset of $\mathbb{R}^n$. Then a search strategy σ is an optimal sequential search procedure if and only if it follows an optimal simultaneous search procedure at every history.*

[7] In the one dimensional case, even if the searcher could conclude search with an option that she had not previously explored, she would never do so. In the worst case, the quality of such an unexplored item would be strictly lower than that of the best previously discovered option. By comparison, allowing the searcher to take an unexplored option can drastically change the solution in Weitzman's model; see Doval (2018).

This needs no separate proof, as the proof of Theorem 1 in Appendix B never makes use of the one-dimensionality of S.

As a consequence, comparative statics with respect to search propensity (Proposition 1, Proposition 2) and on step size after bad news (Proposition 4) readily generalize to the multidimensional case as well, when $Q = Q^{MP}$.

On the other hand, there is no simple analogue to the *most promising area* or *triangulation* when the search space is multidimensional. Moreover, the characterization of optimal search in the case where $Q = Q^{QC}$ does not readily generalize to the multidimensional case; Appendix B.2 highlights some difficulties that arise.

6.3 Algorithm for the Constant Fixed Cost Model

Suppose there are $n \in \mathbb{N}$ objects that can be searched, i.e., $|S| = n$. Let h_t be some history where $t < n$ objects have been discovered already. Let W_t be the search window at this history. Costs are constant and fixed at $c > 0$. Let $M > 0$ be the smallest natural number such that if the searcher searches more than M times after history h_t, her payoff is guaranteed to be less than her payoff to ending search immediately, regardless of the search outcome. Note that M is independent of n.

The following algorithm returns an action at h_t that is part of an optimal sequential search procedure.

1. For each x in S and y in W_t, let $D(x, y) \equiv \max\{d(s, y) | s \in Conv(\{x, y\}),\ q^u_{h_t}(s) \geq z^*_{x_t}\}$.
2. For each y in W_t and $S' \subset S$, let $D(y, S') = \min_{x \in S'} D(x, y)$.
3. For each $i \in \{0, 1, \ldots, M\}$, compute B_i: the subset of W_t of size i that solves:
$$\min_{B \subset W_t, |B|=i} \max_{y \in S/B \cup X_{h_t}} D(y, B \cup X_{h_t}),$$
with $B_0 = \emptyset$.
4. Let B^* be the set $B_i \in \{B_1, \ldots, B_M\}$ that maximizes the searcher's payoff, were he to follow a simultaneous search strategy that searches all items in B_i in some order.
5. The searcher picks some item in B^* to search if it is non-empty, and concludes search otherwise.

Intuitively, the algorithm picks items so that each undiscovered item is as close as possible to some discovered item. That way, by the bounds imposed by Lipschitz continuity, no undiscovered item can be that much better than a discovered item.

More precisely, in the third step, the searcher finds subsets of items (of different cardinalities) that, if searched, would be closest to the remaining undiscovered items. In the fourth step, by the multidimensional analogue of Theorem 1, she picks the subset that maximizes her payoff along the equilibrium path. This payoff can be computed by Lemma 1 as before. Finally, the searcher picks one item in this subset and repeats the procedure until search concludes.

The first step and third step adjust for the right notion of distance. In the case where there is no item or only one item had been discovered so far, $D(y, S')$ is the usual set distance between item y and a set S'. If item x had been discovered and has quality below $z^*_{h_t}$, then we measure the distance between y and the closest point between y and x (not necessarily in S) at which quality could conceivably exceed $z^*_{h_t}$.

The algorithm for multidimensional search reflects Weitzman's intuition (quoted in Section 1) that search in the presence of correlated rewards should favor exploring items most correlated with other unexplored options. Starting at a history where the searcher had not yet explored any items, she plans to explore a set of items that would be closer to all the unexplored items than any other set of the same cardinality. In this prior-free setting, minimizing set distance between explored and unexplored items is analogous to maximizing correlation.

6.3.1 Optimal search set and utility parameters

The only role of costs c and the benefits function U in the algorithm is to determine how many more times the searcher searches along the equilibrium path. Conditional on expecting to search k more items, two searchers will find the same sets of items to be optimal for exploration.

Formally, consider a multidimensional search space model where $Q = Q^{MP}$ and $|S| < \infty$. Suppose there are two searchers, Mary and Bob, with benefits and constant fixed cost of search given by (U, c) and (U', c'), respectively. At a history $h \in H$, suppose σ_M (σ_B) is an optimal deterministic simultaneous search procedure for Mary (Bob). Suppose that starting at history h, search concludes under σ_M and σ_B after the same number of searches. Then σ_B is an optimal simultaneous search procedure for Mike (and σ_M for Bob).

This independence property is helpful in tackling 'robust search space design' problems, considered in Section 6.4.

6.3.2 Complexity Analysis

Note that the first two steps require $O(n^2)$ computations (computing the modified notion of distance for, at most, every pair of items). The fourth and fifth steps collectively require

$O(M)$ computations. The third step requires $O(\binom{n}{1} + \ldots \binom{n}{M}) = O(n^M)$ computations. She repeats these operations at most M times, giving an overall time complexity of $O(Mn^M)$.

The algorithm is therefore polynomial in n, e.g., the number of products that an online marketplace lists.

Indeed, keeping M fixed and letting n grow large is the appropriate limit. A searcher's payoffs will determine the bound M at the very start: she will never make more than M searches. If c is large, or if U_2 is close to zero, then M will be small.

The rational agent paradigm is a dubious benchmark in settings where optimization entails solving NP-hard problems. The polynomial time complexity of optimal multidimensional search in the present model is an encouraging result in this respect.

Additionally, using this algorithm, one can simulate search with learning (e.g., to match moments in data) without compromising on dynamic consistency. Meanwhile, solving analogous Bayesian search models with learning is computationally intractable. Hodgson and Lewis (2020), for example, assume agents only look ahead one period to tackle this issue.

6.4 Search space design

Online marketplaces, like Amazon or Airbnb, routinely make decisions about how to order products. When a shopper searches for a gaming computer, in which order should products be displayed? If he clicks on a black Nike sweatshirt, which related products should be shown?

A constraint is that the user only has so much attention and screen real-estate, so the website cannot display every relevant product at once. It can perhaps show one, or two or four items at a time. It must therefore build out a directed network, linking each product page to a limited number of other products that the buyer can discover next.

How should such a network be constructed? Clearly, ordering matters. Suppose the only way to navigate away from an adult Trek road bike product page is by sampling children's Mongoose bikes. A shopper may quickly get frustrated with the difficulty of finding desirable products and leave without making a purchase. It would also be frustrating if the only related products shown on the product page of a Cuisinart can opener are differently colored but otherwise identical Cuisinart can openers. Consumers may leave for other websites that better expose them to the world of can openers.

Choosing the network of related products can be thought of as choosing the costs of exploration starting at any product. For example, suppose there is a finite set of products in some multidimensional attribute space. Only d related products can be shown on each product page. If B is chosen to be a related product to A, the cost of searching B after A is some $c > 0$ for the searcher. Otherwise, the cost of searching B after A is ∞ for the searcher. A choice of network induces a corresponding search cost function C, which, at any

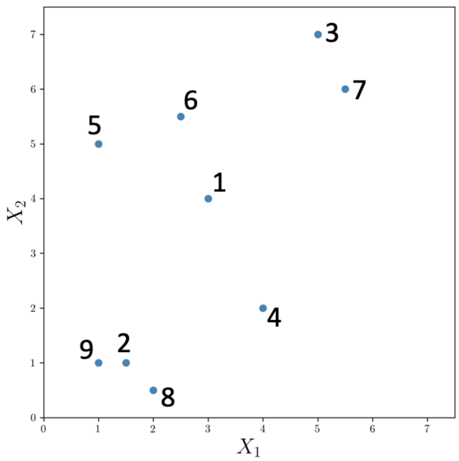

# searches	optimal search set	distance
1	(3, 4)	3.64
2	(3, 4), (1.5, 1)	3.61
3	(3, 4), (1.5, 1), (5, 7)	2.24
4	(3, 4), (1.5, 1), (5, 7), (2, 0.5)	2.24
5	(3, 4), (1.5, 1), (5, 7), (4, 2), (1, 5)	1.58
6	(3, 4), (1.5, 1), (5, 7), (4, 2), (1, 5), (2.5, 5.5)	1.12
7	(3, 4), (1.5, 1), (5, 7), (4, 2), (1, 5), (2.5, 5.5), (5.5, 6)	0.71
8	(3, 4), (1.5, 1), (5, 7), (4, 2), (1, 5), (2.5, 5.5), (5.5, 6), (2.5, 5.5)	0.50
9	(3, 4), (1.5, 1), (5, 7), (4, 2), (1, 5), (2.5, 5.5), (5.5, 6), (2.5, 5.5), (1, 1)	0.00

Figure 4: On the left, the search space consists of 9 points on a plane in $\mathbb{R}^2$. For any U and c, a searcher has an optimal number k of points that she plans to explore along the equilibrium path, starting at the history where she only observes the quality at the point $(3, 4)$. This k (first column of the table) determines the optimal search set (second column) in a constant fixed-costs model. The worst case payoff occurs at an unexplored item that is farthest from the optimal search set (third column). The points are numbered such that any searcher who plans any number of searches can always explore only their ideal optimal set by clicking the next product recommendation, starting with 1 and moving higher. Note that no searcher would ever plan to make two three searches, since they could search twice more and achieve the same worst case outcome. It is therefore of no consequence that the product at $(2, 0.5)$ comes well after fourth place. The ordering favors products that *all* searchers who plan at least four more searches would want to see.

history and for any unexplored product, takes on the value c or ∞.

If there were no constraints on how many related products could be shown, the agent would be able to move from one product to any another. Given the constant fixed costs, it is optimal for her to follow the algorithm from Section 6.3.

This suggests a way for the designer to design the network: choose a network that allows the searcher to discover the set of products she would find under optimal multidimensional spatial search along the equilibrium path (i.e., in the worst case), had there been no constraints on the number of related products that could be displayed. In particular, choose the network so that the searcher need not pass through products that they would rather skip over in an unconstrained search.

Suppose searchers enter the network by first observing some product x_0. How much does the designer need to know about each consumer's c and U to do this?

Surprisingly, it may not need to know anything at all. As Section 6.3.1 elucidates, if the searcher plans to search k more times, the best k points to search, say $X(k) \subset S$, is independent of preferences. The designer simply has to choose a network where for every k, there is a length k path emanating from x_0 that consists of precisely the products in $X(k)$. It is possible to construct such a network if d is sufficiently large. Or to put it differently, for a given d, it is possible to construct such a network if there if the optimal set of k searches has a large enough overlap with the optimal set of $k+1$ searches for several values of k. Figure 4 gives an example where even $d = 1$ suffices.

It is better to dynamically update suggested products if possible: the worst case outcome need not occur, and the searcher may find herself recomputing her optimal exploration set accordingly. The designer can implement the searcher's optimal sequential search procedure by showing an optimal suggested product, if she receives perfect feedback about the searcher's preferences and search outcomes.

Tadelis (2016) argues that such feedback can be useful for search prioritization though we must better understand "how to engineer ways in which more accurate feedback is generated". If the designer's information is coarser (i.e., she infers the size and direction of news from consumer actions like posting questions, clicking on reviews, adding items to the basket etc.), she can try to let the searcher find her optimal next product by showing an array of suggested products. To this end, increasing d is useful when the designer is more uncertain about how a searcher perceives the current product.

7 Conclusion

This paper puts forth a simple and tractable model of costly spatial search with forward looking agents. One objective is to show that model can produce dynamics that match those found in the analysis of consumer search data. Moreover, optimal search here is shown to be computationally tractable, which makes it both a more plausible model of economic agents and a simpler rule to simulate. These properties may make the model amenable to empirical applications.

To that end, it appears feasible to push the model further. For example, an item's observable attributes may not pin down its quality. Indeed there may be several different items with the same observable qualities. This suggests modeling the mapping from attributes to qualities as a correspondence rather than a function as done here.

Similarly, qualities may not be perfectly observable as assumed in our setting. Instead, the searcher may observe only a signal in the vicinity of any explored item's true quality. Our model can easily be adapted to address the case where this signal is adversarially chosen. It would be interesting to extend results to other kinds of information structures.

References

Adam, Klaus, "Learning while searching for the best alternative," *Journal of Economic Theory*, 2001, *101* (1), 252–280.

Aghion, Philippe, Patrick Bolton, Christopher Harris, and Bruno Jullien, "Optimal learning by experimentation," *The review of economic studies*, 1991, *58* (4), 621–654.

Banerjee, Abhijit V, Sylvain Chassang, and Erik Snowberg, "Decision theoretic approaches to experiment design and external validity," in "Handbook of Economic Field Experiments," Vol. 1, Elsevier, 2017, pp. 141–174.

Bikhchandani, Sushil and Sunil Sharma, "Optimal search with learning," *Journal of Economic Dynamics and Control*, 1996, *20* (1-3), 333–359.

Blake, Thomas, Chris Nosko, and Steven Tadelis, "Returns to consumer search: Evidence from ebay," in "Proceedings of the 2016 ACM Conference on Economics and Computation" 2016, pp. 531–545.

Bronnenberg, Bart J, Jun B Kim, and Carl F Mela, "Zooming in on choice: How do consumers search for cameras online?," *Marketing science*, 2016, *35* (5), 693–712.

Callander, Steven, "Searching and learning by trial and error," *American Economic Review*, 2011, *101* (6), 2277–2308.

Carroll, Gabriel, "Robustness in mechanism design and contracting," *Annual Review of Economics*, 2019, *11*, 139–166.

Chassang, Sylvain, "Mostly prior-free asset allocation," *Princeton University William S. Dietrich II Economic Theory Center Research Paper*, 2016, (077_2016).

Doval, Laura, "Whether or not to open Pandora's box," *Journal of Economic Theory*, 2018, *175*, 127–158.

Garfagnini, Umberto and Bruno Strulovici, "Social experimentation with interdependent and expanding technologies," *Review of Economic Studies*, 2016, *83* (4), 1579–1613.

Gittins, John, "A dynamic allocation index for the sequential design of experiments," *Progress in statistics*, 1974, pp. 241–266.

Guo, Yingni and Eran Shmaya, "Robust monopoly regulation," *arXiv preprint arXiv:1910.04260*, 2019.

Hansen, Pierre, Brigitte Jaumard, and Shi-Hui Lu, “Global optimization of univariate Lipschitz functions: I. Survey and properties,” *Mathematical programming*, 1992, *55* (1), 251–272.

Hodgson, Charles and Gregory Lewis, “You can lead a horse to water: Spatial learning and path dependence in consumer search,” 2020.

Hörner, Johannes and Andrzej Skrzypacz, “Learning, experimentation and information design,” 2017.

Kreps, David M, “Decision problems with expected utility critera, i: upper and lower convergent utility,” *Mathematics of Operations Research*, 1977, *2* (1), 45–53.

Lazear, Edward P, “Economic imperialism,” *the Quarterly Journal of economics*, 2000, *115* (1), 99–146.

Malladi, Suraj, “Delegated Costly Screening,” 2021.

Manski, Charles F, “Minimax-regret treatment choice with missing outcome data,” *Journal of Econometrics*, 2007, *139* (1), 105–115.

McCall, John Joseph, “Economics of information and job search,” *The Quarterly Journal of Economics*, 1970, pp. 113–126.

Rothschild, Michael, “Searching for the lowest price when the distribution of prices is unknown,” in “Uncertainty in Economics,” Elsevier, 1978, pp. 425–454.

Schlag, Karl H and Andriy Zapechelnyuk, “Robust sequential search,” *Theoretical Economics*, 2021, *16* (4), 1431–1470.

Shubert, Bruno O, “A sequential method seeking the global maximum of a function,” *SIAM Journal on Numerical Analysis*, 1972, *9* (3), 379–388.

Slivkins, Aleksandrs, “Introduction to multi-armed bandits,” *arXiv preprint arXiv:1904.07272*, 2019.

Stigler, George J, “The economics of information,” *Journal of political economy*, 1961, *69* (3), 213–225.

Sukharev, Aleksei G, “Best sequential search strategies for finding an extremum,” *USSR Computational Mathematics and Mathematical Physics*, 1972, *12* (1), 39–59.

Tadelis, Steven, “Reputation and feedback systems in online platform markets,” *Annual Review of Economics*, 2016, *8*, 321–340.

Tetenov, Aleksey, "Statistical treatment choice based on asymmetric minimax regret criteria," *Journal of Econometrics*, 2012, *166* (1), 157–165.

Timonov, LN, "Algorithm for search of a global extremum," *Engineering Cybernetics*, 1977, *15* (3), 38–44.

Urgun, Can and Leeat Yariv, "Retrospective search: Exploration and ambition on uncharted terrain," Technical Report, National Bureau of Economic Research 2021.

Weitzman, Martin L, "Optimal search for the best alternative," *Econometrica: Journal of the Econometric Society*, 1979, pp. 641–654.

Wong, Yu Fu, "Spatial experimentation," *Available at SSRN 3611751*, 2021.

A Optimal Sequential Search Algorithm: Constant Fixed Costs Model

The boundary points of W_{h_t} and the technologies in $X_{h_t} \cap W_{h_t}$ partition W_{h_t} into sub-intervals. Let $mesh(W_{h_t}, X_{h_t})$ be the length of the longest of these sub-intervals. If $0 \in W_{h_t}/X_{h_t}$, let l_{h_t} denote the length of the sub-interval in W_{h_t} containing 0 ; otherwise, let $l_{h_t} = 0$. Likewise, if $1 \in W_{h_t}/X_{h_t}$, let r_{h_t} denote the length of the sub-interval in W_{h_t} containing 1; otherwise, let $r_{h_t} = 0$. The *weighted mesh at* h_t is $w - mesh(W_{h_t}, X_{h_t}) \equiv \max\{mesh(W_{h_t}, X_{h_t}),\ 2l_{h_t},\ 2r_{h_t}\}$.

For any $k \in \mathbb{N}$, let $T(k, h_t)$ and denote a solution to the following problem:

$$\min_{T \in S/X_{h_t}; |T|=k} w - mesh(W_{h_t}, T \cup X_{h_t}).$$

Label the elements of $T(k, h_t)$ as $x_{t+1}, \ldots, x_{t+k}$ in any order. Let

$$h_{t+k} = \{(x_0, z_0), \ldots, (x_t, z_t), (x_{t+1}, q^u_{h_t}(x_{t+1})), \ldots, (x_{t+k}, q^u_{h_t}(x_{t+k}))\},$$

and let $v(k, h_t) \equiv p(h_{t+k}, q^u_{h_{t+k}})$.

Proposition **6.** *Let $Q = Q^{QC}$ or $Q = Q^{MP}$, and consider a constant fixed cost model. At every history $h_t \in H$:*

1. *Let $D_{h_t} = p(h_t, q^u_{h_t}) - p(h_t)$.*
2. *Let $n = \lfloor \frac{D_{h_t}}{c} \rfloor$.*
3. *Let k^* be a solution to $\max_{k=0,\ldots,n} v(k, h_t)$.*
4. *Let $\sigma(h_t) \in T(k^*, h_t)$.*

Then σ is an optimal sequential search procedure.

We omit the proof since it follows directly from Theorem 1, Theorem 2 and Lemma 1.

B Proofs

B.1 Proof of Theorem 1

Let $h_t = \{(x_i, z_i)\}_{i=0}^{t-1} \in H$. For each $x_i \in X_{h_t}$ and $y \in S$, let $f_{h_t,x_i}(y) = L||y - x_i|| + z_i$. Let $f_{h_t} = \min_{x \in X_{h_t}} f_{h_t,x}$. Recall that $q^u_{h_t}$ is the upper envelope of the quality indices in Q_{h_t} for any history $h_t \in H$.

Lemma **2.** *Suppose that $Q = Q^{MP}$. Then $q^u_{h_t} \in Q_{h_t}$ for any $h_t \in H$. Moreover, $q^u_{h_t} = \min\{f_{h_t}, 1\}$.*

Proof of Lemma 2. We proceed by proving three claims.

Claim 1: $\min\{f_{h_t}, 1\}$ is L-Lipschitz.

Let $x, y \in S$. Then there exists some $x_i, x_j \in X_{h_t}$ such that $|f_{h_t}(x) - f_{h_t}(y)| = |f_{h_t,x_i}(x) - f_{h_t,x_j}(y)|$. Suppose without loss of generality that $f_{h_t,x_i}(x) \geq f_{h_t,x_j}(y)$. Then

$$\begin{aligned}
|f_{h_t}(x) - f_{h_t}(y)| &= f_{h_t,x_i}(x) - f_{h_t,x_j}(y) \\
&\leq f_{h_t,x_j}(x) - f_{h_t,x_j}(y) \\
&= L||x - x_j|| + z_j - L||y - x_j|| - z_j \\
&= L||x - x_j|| - L||y - x_j|| \\
&\leq L||x - y||,
\end{aligned}$$

where the first inequality follows from the definition of f_{h_t}. Therefore, f_{h_t}, and therefore $\min\{f_{h_t}, 1\}$, is L-Lipschitz.

Claim 2: Every quality index in Q_{h_t} is bounded above pointwise by $\min\{f_{h_t}, 1\}$.

Let $h_t = \{(x_i, z_i)\}_{i=0}^{t-1} \in H$, and let $q' \in Q_{h_t}$. For any $x_i \in T$ and $y \in S$, $|q'(y) - q'(x_i)| \leq L||y - x_i||$; Moreover, note that $|f_{h_t,x_i}(y) - f_{h_t,x_i}(x_i)| = f_{h_t,x_i}(y) - q'(x_i) = L||y - x_i||$. Together, this implies that $q'(y) \leq f_{h_t,x_i}(y)$ for all $x_i \in T$. Therefore, $q'(y) \leq \min\{f_{h_t}(y), 1\}$ for all $q' \in Q_{h_t}$.

Claim 3: $\min\{f_{h_t}, 1\}$ is consistent at h_t.

For any $x_i \in X_{h_t}$ and $q' \in Q_{h_t}$, $z_i = q'(x_i) \leq f_{h_t}(x_i) \leq f_{h_t,x_i}(x_i) = z_i$, where the first inequality follows from *Claim 2*, and the second is by the definition of f_{h_t}. Therefore $f_{h_t}(x_i) = \min\{f_{h_t}(x_i), 1\} = z_i$ for all $x_i \in X_{h_t}$, so $\min\{f_{h_t}, 1\}$ is consistent.

The first and third claims imply that $\min\{f_{h_t}, 1\} \in Q_{h_t}$. So by the second claim, f_{h_t} is the upper envelope of the quality indices in Q_{h_t}, i.e., $q^u_{h_t} = \min\{f_{h_t}, 1\}$. □

Let $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1}$ and $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1}$ be two partial histories that differ only by quality. Suppose $z'_i \geq z''_i$ for all i. Then we say h'_t *dominates* h''_t in *quality*.

Lemma **3.** *Suppose that $Q = Q^{MP}$. If h'_t dominates h''_t in quality, then*

$$\max_{q' \in Q_{h'_t}, x \in S} q'(x) \geq \max_{q' \in Q_{h''_t}, x \in S} q'(x).$$

Proof of Lemma 3. The statement of the lemma is equivalent to showing that:

$$\max_{x \in S} q^u_{h'_t} \geq \max_{x \in S} q^u_{h''_t}.$$

Let $T \equiv X_{h'_t} = X_{h''_t}$ be the set of searched technologies in $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1}$ and $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1}$. By Lemma 2, it suffices to show that $f_{h'_t,x_i}(y) = ||y - x_i|| + z'_i \geq ||y - x_i|| + z''_i = f_{h''_t,x_i}$ for every $x_i \in T$ and $y \in S$. This follows immediately from the assumption that $z'_i \geq z''_i$ for all i. $\square$

Lemma 4. *Suppose that $Q = Q^{MP}$. Fix some history $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1} \in H$ and let x be an optimal technology at h'_t. Let $T = X^*_{h'_t} \setminus \{x\}$. For each $x_i \in T$, let $\epsilon_i \in \mathbb{R}$. Consider an alternate history $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1} \in H$ where $z''_i = z'_i + \epsilon_i$ for $x_i \in T$ and $z''_i = z'_i$ otherwise. Suppose moreover that $Q_{h'_t}$ and $Q_{h''_t}$ are nonempty . The searcher's payoff to concluding search in h''_t is greater than or equal to the searcher's payoff to concluding search at h'_t.*

Proof of Lemma 4. If T is empty, the statement is trivially true. Suppose T is non-empty.

Case 1: If $\epsilon_i \leq 0$ for each $x_i \in T$, then note that the quality of the searcher's chosen technology is the same in both h''_t and h'_t, but by Lemma 3, the adversarially selected quality of the best technology in S is weakly higher in h'_t. Therefore, the searcher's benefit to concluding search is weakly greater at h''_t than at h'_t.

Case 2: Suppose there exists some $x_i \in T$ such that $\epsilon_i > 0$. Let ϵ denote the largest ϵ_i among all i such that $x_i \in T$. Consider a third partial history $h'''_t = \{(x_i, z'_i + \epsilon)\}_{i=0}^{t-1}$, i.e., where the quality of *all* technologies searched in h'_t is higher than in h'_t by ϵ. It is immediate by definition that $q^u_{h'''_t} \leq q^u_{h'_t} + \epsilon$, with equality if the upper bound of 1 does not bind, so

$$\max_{q' \in Q_{h'''_t}, x \in S} q'(x) \leq \epsilon + \max_{q' \in Q_{h'_t}, x \in S} q'(x).$$

But then the searcher's benefit to concluding search is weakly higher at h'''_t than at h'_t, because by assumption, $U(a + \epsilon, b + \epsilon) \geq U(a, b)$ for any $a, b \in \mathbb{R}_+$.

The quality of the searcher's chosen technology is the same in both h''_t and h'''_t. But by Lemma 3 again, the searcher's benefit to concluding search is weakly greater at h''_t than at h'''_t, and therefore weakly greater at h''_t than at h'_t.

Recall C is independent of the qualities of searched technologies. Therefore, the searcher's expected costs of experimentation would be the same upon concluding search after partial h'_t or h''_t. Therefore, the searcher's payoff is weakly greater if she concludes search at h''_t than at h'_t. $\square$

Let $q^A_{h_t} = \min\{q^u_{h_t}, q^u_{h_t}(z^*_{h_t})\}$. For all $h_t \in H$, let $\sigma^A_d(h_t, x) = q^A_{h_t}$ for all $x \neq \emptyset$, and $\sigma^A_d(h_t, \emptyset) = q^u_{h_t}$.

Lemma 5. *Suppose that $Q = Q^{MP}$. Then σ^A_d is a feasible strategy for the adversary.*

Proof of Lemma 5. By Lemma 2, it suffices to prove that $q^A_{h_t} \in Q_{h_t}$. Clearly, $q^A_{h_t}$ is L-Lipschitz since it is the minimum of an L-Lipschitz function and a constant. By definition, it is also consistent. □

LEMMA **6.** *Suppose that $Q = Q^{MP}$ and let $\sigma \in \Gamma_{h_t}$. Then σ is an optimal simultaneous search procedure at h_t if and only if σ is a best-response to σ^A_d in the simultaneous search game at $h_t \in H$.*

Proof of Lemma 6. Suppose first that σ is an optimal simultaneous search procedure at h_t. Let T denote the set of technologies explored under σ after h_t. Because $\sigma \in \Gamma_{h_t}$, T is independent of the adversary's strategy. Then by Lemma 4, σ^A_d weakly dominates any other strategy for the adversary.

Therefore, if σ is a best response to σ^A_d, then (σ, σ^A_d) is a Nash equilibrium of the simultaneous search game at h_t, i.e., σ is an optimal simultaneous search procedure.

For the other direction, suppose for contradiction that σ is not a best response to σ^A_d but that (σ, σ^A) is a Nash equilibrium of the simultaneous search game at h_t. Let σ' be a best response to σ^A_d, so that (σ', σ^A_d) is another equilibrium. Let $\preceq$ represent the searcher's preferences over outcomes.

First, $(\sigma', \sigma^A) \preceq (\sigma, \sigma^A)$, by the assumption that the latter is a Nash equilibrium.

Next, $(\sigma, \sigma^A) \preceq (\sigma, \sigma^A_d)$: the former is an equilibrium of the zero-sum game, so adversary is weakly worse off if she deviates, and the searcher is weakly better off.

However, it is also true that $(\sigma, \sigma^A_d) \preceq (\sigma, \sigma^A)$ because σ^A_d weakly dominates any other strategy for the adversary, leaving the searcher with a weakly smaller payoff.

Finally, $(\sigma, \sigma^A_d) \prec (\sigma', \sigma^A_d)$, since σ is not a best response to σ^A_d while σ' is a best response.

Putting this all together, $(\sigma', \sigma^A) \preceq (\sigma, \sigma^A) \sim (\sigma, \sigma^A_d) \prec (\sigma', \sigma^A_d)$, i.e., $(\sigma', \sigma^A) \prec (\sigma', \sigma^A_d)$. In other words, when the searcher plays σ', she is strictly better off when the adversary plays σ^A_d over σ^A. This implies the adversary is strictly worse off playing σ^A_d over σ^A when the searcher plays σ', contradicting the fact that σ^A_d weakly dominates any other strategy for the adversary.

Therefore if σ is not a best response to σ^A_d, it is not an optimal simultaneous search procedure. □

LEMMA **7.** *Suppose that $Q = Q^{MP}$ and that σ follows an optimal simultaneous search procedure at every history. Let $h_t \in H$, and let $\sigma' \in \Gamma_{h_t}$ be a simultaneous search procedure at h_t which replicates the searcher's path of play under (σ, σ^A_d), conditional on reaching h_t. Then σ' is an optimal simultaneous search procedure at h_t.*

Proof of Lemma 7. Fixing the adversary's strategy as σ_d^A, Lemma 6 implies that σ is a conserving strategy for the searcher, i.e., the highest payoff possible at h_t equals the highest payoff possible after the searcher takes the action $\sigma(h_t)$.

Next, let h be an infinite history where the searcher never concludes search, and consider the best payoff possible for the searcher at an alternate history that matches h for the first t periods: $\sup_{h' \in H, h'(t)=h(t)} \min_{q \in Q_{h'}} p(h', q)$. Because C is bounded away from zero, the searcher's best achievable payoff decreases without bound as she searches indefinitely, i.e., $\sup_{h' \in H, h'(t)=h(t)} \min_{q \in Q_{h'}} p(h', q) \to -\infty = p(h, \cdot)$ as $t \to \infty$.

Therefore conserving strategies are optimal in this setting (e.g., see Kreps (1977)), i.e., σ is an optimal strategy for the searcher when the adversary plays σ_d^A.

This implies that σ' is a best response to σ_d^A in the simultaneous search game, so by Lemma 6, σ' is an optimal simultaneous search procedure. □

LEMMA **8.** *Suppose that $Q = Q^{MP}$. If σ follows an optimal simultaneous search procedure at every history, then (σ, σ_d^A) is a sub-game perfect equilibrium.*

Proof of Lemma 8. First we show that the searcher's strategy is unimprovable. Let $h_t \in H$. Let σ' denote some one-shot deviation from σ at h_t. Let σ'' and $\sigma''' \in \Gamma_{h_t}$ be simultaneous search procedures at h_t that replicate the searcher's path of play under (σ, σ_d^A) and (σ', σ_d^A), respectively.

Then σ' is a strict improvement over σ only if σ''' is a strict improvement over σ'' in the simultaneous search game at h_t. But by Lemma 7, σ'' is an optimal simultaneous search procedure and therefore a best response to σ_d^A in Γ_{h_t}. Therefore, σ''' is not a strict improvement over σ'', and so σ' is not a strict improvement over σ.

Next, we prove that σ_d^A is unimprovable. Fix the strategy of the searcher to be σ and suppose the searcher picks technology x_t at history h_t. Suppose that the adversary deviates at this history to $q' \neq q_{h_t}^A$ and returns to following σ^A thereafter. Denote this history as h'_{t+1} There are three cases to consider.

Case 1: $q'(x_t) = q_{h_t}^A(x_t)$. In this case, all future histories would proceed as if there was no deviation at all, so this is not a strict improvement for the adversary.

Case 2: $q'(x_t) > q_{h_t}^A(x_t)$. Suppose for a moment that the searcher behaves as if the quality of x_t was $q_{h_t}^A(x_t)$ and continues to follow at h'_{t+1} what was the optimal simultaneous search procedure at h_t, i.e., σ^*_{s,h_t}. The searcher would be better off upon concluding search, by Lemma 4, than she would have been had the adversary not deviated from σ^A at h_t. If the searcher plays σ and follows $\sigma^*_{s,h'_{t+1}}$ at h'_{t+1}, she is better off still. Therefore the searcher is weakly better off when the adversary makes this one-shot deviation.

Case 3: $q'(x_t) < q_{h_t}^A(x_t)$. As in the previous case, we first consider what would happen if the searcher followed σ^*_{s,h_t} at h'_{t+1} onward. Clearly $q_{h'_{t+1}} \leq q_{h_t}$, and the quality of the best searched technology is the same whether the adversary makes this one-shot deviation or not. By Lemma 3, if the searcher follows σ^*_{s,h_t} at h'_{t+1} onward, she is weakly better off when the adversary deviates at h_t. This implies that even when she follows σ, she is weakly better off when the adversary deviates, i.e., the adversary is weakly worse off.

Since the adversary has no strictly profitable one-shot deviation, the strategies (σ, σ_d^A) constitute a sub-game perfect equilibrium. □

LEMMA **9.** *Suppose that $Q = Q^{MP}$. If σ is an optimal sequential search procedure then σ follows an optimal simultaneous search procedure at every history partial $h_t \in H$.*

Proof of Lemma 9. Since σ is an optimal sequential procedure, there exists some strategy of the adversary σ^A such that (σ, σ^A) is a sub-game perfect equilibrium. In particular, let $h_t \in H$; then (σ, σ^A) is a Nash equilibrium of the sub-game starting at h_t. Let σ' be some strategy that follows an optimal simultaneous search procedure at every history. Then (σ', σ_d^A) is also a Nash equilibrium of the sub-game starting at h_t by Lemma 8.

Note that because (σ', σ_d^A) is an equilibrium of the zero-sum sub-game at h_t, the adversary is weakly better off at (σ, σ_d^A) as any deviation leaves the searcher weakly worse-off. And over (σ, σ_d^A), the adversary is weakly better off at (σ, σ^A), since σ^A is a best response to σ by assumption. And again because (σ, σ^A) is an equilibrium, the adversary's payoff weakly improves under the strategy profile (σ', σ^A). And finally, the adversary's payoff weakly improves from (σ', σ^A) if she best responds to the searcher instead at $(\sigma', \sigma^A)_d$.

This loop of weak inequalities implies that the adversary's payoffs are identical at all of these strategy profiles. In particular, this means that (σ, σ_d^A) is an equilibrium in the sub-game at h_t. Therefore a strategy $\sigma'' \in \Gamma_{h_t}$ that replicates the searcher's path of play under (σ, σ_d^A) is an optimal simultaneous search procedure by Lemma 6.

In other words, σ follows an optimal simultaneous search procedure at h_t. Since the choice of h_t was arbitrary, the result follows. □

Proof of Theorem 1. Lemma 8 and Lemma 9 together give the result. □

B.2 Proof of Theorem 2

The main difference between the proofs of Theorem 1 and Theorem 2 is that analogues of Lemma 2 and Lemma 3 no longer holds when $Q = Q^{QC}$.

To see this, consider the following counter-example: Let $S = [0, 4]$ and $L = 1$. Denote by h_3 the partial history where technologies $\{0, 2, 3, 4\}$ have been searched and all have quality equal to 0, i.e., $h_3 = \{(0, 0), (2, 0), (3, 0), (4, 0)\}$.

First, note that the upper envelope of Q_{h_3} is a saw-tooth shaped function and therefore not quasiconcave.

Next, note that the highest possible quality for some technology under some $q \in Q_{h_3}$ is equal to 1. This is uniquely achieved at:

$$q(x) = \begin{cases} x & 0 \leq x < 1 \\ 2 - x & 1 \leq x < 2 \\ 0 & 2 \leq x \leq 4. \end{cases}$$

Now consider the history $h_3' = \{(0,0),(2,0),(3,0.5),(4,0)\}$, which dominates h_3 in quality. Since every quality index in $Q_{h_3'}$ is quasiconcave, it must now be the case that $q'(1) = 0$ for every $q' \in Q_{h_3}$. The highest possible quality for some technology under some $q' \in Q_{h_3'}$ is equal to 0.75. This is achieved at:

$$q'(x) = \begin{cases} 0 & 0 \leq x < 2 \\ x - 2 & 2 \leq x < 2.75 \\ 3.5 - x & 2.75 \leq x < 3.5 \\ 0 & 3.5 \leq x \leq 4. \end{cases}$$

However, an analogue of Lemma 5 and a weaker form of Lemma 3 hold, which suffice for the proof of Theorem 2. To this end, let $q^u_{h_t}$ again denote the upper envelope of quality indices in Q_{h_t} for any $h_t \in H$, and let $q^A_{h_t} = \min\{q^u_{h_t}, z^*_{h_t}\}$.

Lemma **10.** *Suppose that $Q = Q^{QC}$. Then for all $h_t \in H$, $q^A_{h_t} \in Q_{h_t}$*

Proof. The argument that $q^A_{h_t}$ is L-Lipschitz and consistent is exactly as in the proofs of Lemma 2 and Lemma 5. It only remains to be shown that $q^A_{h_t}$ is quasiconcave.

Let $h_t = \{x_i, z_i\}_{i=1}^{t-1} \in H$. If $q' \in Q_{h_t}$, then by quasiconcavity, q' is non-decreasing on $[0, \min X^*_{h_t})$ and non-increasing on $(\max X^*_{h_t}, 1]$. This implies $q^u_{h_t}$ is non-decreasing on $[0, \min X^*_{h_t})$ and non-increasing on $(\max X^*_{h_t}, 1]$. Therefore $q^A_{h_t}$ is quasiconcave. $\square$

Lemma **11.** *Suppose that $Q = Q^{QC}$. If $h_t' = \{(x_i, z_i')\}_{i=0}^{t-1}$ dominates $h_t'' = \{(x_i, z_i'')\}_{i=0}^{t-1}$ in quality, and $z^*_{h_t'} = z^*_{h_t''}$, then*

$$\max_{q' \in Q_{h_t'}, x \in S} q'(x) \geq \max_{q' \in Q_{h_t''}, x \in S} q'(x).$$

Proof of Lemma 11. First we argue that $q^A_{h_t'} \geq q^A_{h_t''}$.

Since h_t' dominates h_t'' and $z^*_{h_t'} = z^*_{h_t''}$, it follows that $X^*_{h_t''} \subset X^*_{h_t'}$. Therefore, $q^A_{h_t'} \geq q^A_{h_t''}$ on $[\min X^*_{h_t'}, \max X^*_{h_t'}]$.

Next, it follows from Lemma 10 that $q^A_{h'_t}$ and $q^A_{h''_t}$ are non-decreasing on $[0, \min X^*_{h'_t})$. This implies that $\max\{q^A_{h'_t}, q^A_{h''_t}\}$ is also non-decreasing on this interval. Reasoning analogously, $\max\{q^A_{h'_t}, q^A_{h''_t}\}$ is non-increasing on $(\min X^*_{h'_t}, 0]$. Finally, observe that $\max\{q^A_{h'_t}, q^A_{h''_t}\}$ is consistent with what the searcher had observed so far at h'_t, i.e., $\max\{q^A_{h'_t}, q^A_{h''_t}\} \in Q_{h'_t}$. Therefore, $q^A_{h'_t} = \max\{q^A_{h'_t}, q^A_{h''_t}\} \geq q^A_{h''_t}$.

Now let $q'' \in Q_{h''_t}$. It suffices to show that there exists a $q' \in Q_{h'_t}$ such that $\max_{x \in S} q'(x) = \max_{x \in S} q''(x)$.

Define q' as follows: $q'(x) = q^A_{h'_t}(x)$ if $q''(x) \leq z^*_{h'_t}$ and $q'(x) = q''(x)$ otherwise. By the preceding, $q^A_{h'_t} \geq q^A_{h''_t} \geq \max\{q'', z^*_{h''_t}\}$. Moreover, because $\{x \in [0,1] | q''(x) > q^A_{h'_t}(x)\} \subset [\min X^*_{h'_t}, \max X^*_{h'_t}]$, it is clear that q' is quasiconcave and that $q' \in Q_{h'_t}$.

Since $\max_{x \in S} q'(x) = \max_{x \in S} q''(x)$, the result follows. ☐

Proof of Theorem 2. The proofs of the analogous lemmas to those in Appendix B.1 are identical, with Lemma 11 in place of Lemma 3 whenever the latter is referenced. ☐

B.3 Remaining Proofs from Section 4

Proof of Lemma 1. See the proof of Lemma 6. The proof when $Q = Q^{QC}$ is identical. ☐

Proof of Proposition 1. First we construct $\underline{L}$. If $L = 0$ and Q_h is nonempty, then clearly there is no value in search, as Q_h is a singleton containing only a constant function. Recall that C is bounded away from 0, so suppose that $c > 0$ is such that $C(x, h') > c$ for all $x \in S$ and $h' \in H$. Let $\epsilon > 0$ be small enough so that $U(z^*_h, z^*_h) - U(z^*_h, z^*_h + \epsilon) < c$. By compactness of S, there exists $\underline{L} > 0$ small enough so that $q^u_h < z^*_h + \epsilon$ when search complexity L is such that $L \leq \underline{L}$ and Q_h is nonempty. For contradiction, suppose search does not conclude immediately when $L \leq \underline{L}$ and Q_h is nonempty: there is some history h' after h on the equilibrium path at which the searcher concludes search. Then by Theorem 1 or Theorem 2 and by Lemma 1, the searcher's anticipated payoff is

$$
\begin{aligned}
p(h', q^u_{h'}) &= U(z^*_h, \max_{x \in S} q^u_{h'}(x)) - \sum_{i=1}^{t} C(x_i, h'(i)) \\
&\leq U(z^*_h, z^*_h) - c \\
&< U(z^*_h, z^*_h + \epsilon) \\
&< p(h, \max_{x \in S} q^u_h(x)),
\end{aligned}
$$

a contradiction. Therefore concluding search immediately is optimal if $L \leq \underline{L}$.

Next we construct $\overline{L}$. Let $\delta = U(z^*_h, z^*_h) - U(z^*_h, 1)$. Let $n = \lceil \frac{\delta}{c} \rceil$. Let h' be any history on the equilibrium path at which there have been at least n searches, i.e., $t \geq n$. Then

$$\begin{aligned} p(h', q^u_{h'}) &= U(z^*_h, \max_{x \in S} q^u_{h'}(x)) - \sum_{i=1}^{t} C(x_i, h'(i)) \\ &\leq U(z^*_h, z^*_h) - n \cdot c \\ &< U(z^*_h, 1). \end{aligned}$$

Therefore, it suffices to construct $\overline{L}$ such that if $L \geq \overline{L}$, then after any history with n searches where the adversary plays σ^A_d, $\max_{x \in S} q^u_{h'}(x) = 1$. This would imply that the searcher is better off concluding search immediately.

Pick $n+1$ points $x'_1, \ldots, x'_{n+1} \in S/X_h$; let T denote the union of these points and X_h and let d be the minimum over distances between two points in T. Let $\overline{L} \geq \frac{1-z^*_h}{d}$ and suppose $L > \overline{L}$. Consider a history h' after h where some $x_1, \ldots, x_n \in S/X_h$ are searched in some order, with the quality of x_i being $q^A_h(x_i)$. There exists by construction a $q \in Q_{h'}$ such that $q(x_i) = q^A_h(x_i)$ for all i, but $q(x'_j) = 1$ for some j. Therefore $\max_{x \in S} q^u_{h'}(x) = 1$ for any history after which there are n searches. $\square$

When considering comparative statics with respect to different levels of search complexity, say L' and L'', we subscript variables by L' or L'' to denote which case we are considering.

Lemma **12.** *Let $Q = Q^{MP}$ or $Q = Q^{QC}$. Let $0 < L' < L''$ be two different levels of search complexity. Let $h = \{(x_i, z_i)\}_{i=0}^{t} \in H$ be such that $z_i = z_j = z$ for all $i, j \in \{0, \ldots, t\}$, and let $y \in S/X_h$. Let $h' \in H$ be the history immediately after at which y is searched and has quality z. Then,*

$$\max_{x \in S} q^u_{h,L'}(x) - \max_{x \in S} q^u_{h',L'}(x) \leq \max_{x \in S} q^u_{h,L''}(x) - \max_{x \in S} q^u_{h',L''}(x).$$

Moreover, if $\max_{x \in S} q^u_{h,L'}(x) - \max_{x \in S} q^u_{h',L'}(x) > 0$, the preceding inequality is strict.

Proof. Note that because $Q_{h'} \subset Q_h$, $\max_{x \in S} q^u_{h,L''}(x) - \max_{x \in S} q^u_{h',L''}(x) \geq 0$. So, the result holds when $\max_{x \in S} q^u_{h,L'}(x) - \max_{x \in S} q^u_{h',L'}(x) = 0$.

Suppose that $\max_{x \in S} q^u_{h,L'}(x) - \max_{x \in S} q^u_{h',L'}(x) > 0$.

Let $\underline{x}, \overline{x} \in X_h$ which are the closest previously searched technologies to the left and right of y (and $\min S$ or $\max S$, respectively, if there are no such technologies). Define $\underline{x_s}, \overline{x_s}' \in X_h$ similarly as the endpoints in $X_h \cup \{\min S, \max S\}$ of the sub-interval containing the second largest peak of q^u_h.

Let $f(L) \equiv \max_{x \in [\underline{x}, \overline{x}]} q(x)$. Similarly, let $g(L) \equiv \max_{x \in [\underline{x_s}, \overline{x_s}]} q(x)$.

It is readily verified (for example, by Lemma 2 and an analogous result for the $Q = Q^{QC}$ case) that $f(L) = z + D(\frac{\overline{x}, \underline{x}}{2}) \cdot L$, where $D(a, b) \equiv \frac{b-a}{2}$ if $a, b \in X_h$, and $D(a, b) \equiv b - a$ otherwise. Similarly, let $g(L) = z + D(\frac{\overline{x_s}, \underline{x_s}}{2}) \cdot L$.

For the remainder of the proof, we consider only the case where $\underline{x}, \overline{x}, \underline{x_s}, \overline{x_s} \in X_h$. We obtain the same conclusion when one or more of $\underline{x}, \overline{x}, \underline{x_s}, \overline{x_s}$ are not in X_h. There are two cases to consider.

Case 1: Suppose that $\max\{y-\underline{x}, \overline{x}-y\} < \overline{x_s}-\underline{x_s}$. This implies that at h', $\max_{x\in S} q^u_{h',L'}(x) = g(L')$ and $\max_{x\in S} q^u_{h',L''}(x) = g(L'')$.

Now $f(L) - g(L) = L \cdot \frac{\overline{x}-\underline{x}-\overline{x_s}+\underline{x_s}}{2} > 0$ is linear in L with a positive slope, which implies $f(L) - g(L)$ is strictly increasing in L. Therefore $f(L') - g(L') \leq f(L'') - g(L'')$, which is the desired result.

Case 2: Suppose that $\max\{y - \underline{x}, \overline{x} - y\} \geq \overline{x_s} - \underline{x_s}$. Suppose without loss of generality that $y - \underline{x} \geq \overline{x} - y$.

Then at history h',

$$\max_{x\in S} q^u_{h',L'}(x) = \max_{x\in S\cap[\underline{x},y]} q^u_{h',L'}(x),$$

and

$$\max_{x\in S} q^u_{h',L''}(x) = \max_{x\in S\cap[\underline{x},y]} q^u_{h',L''}(x).$$

Let $\alpha = \frac{y-\underline{x}}{\overline{x}-\underline{x}}$. Then by the property of similar triangles,

$$\max_{x\in S} q^u_{h',L'}(x) = \alpha \max_{x\in S} q^u_{h,L'}(x),$$

and

$$\max_{x\in S} q^u_{h',L''}(x) = \alpha \max_{x\in S} q^u_{h,L''}(x).$$

Since $(1-\alpha)\max_{x\in S} q^u_{h',L'}(x) < (1-\alpha)\max_{x\in S} q^u_{h',L''}(x)$, the result follows. □

Proof of Proposition 2. Suppose for contradiction that there is an optimal sequential search procedure that does not conclude search immediately when complexity is L'. Let h' be the history at which search ends along the equilibrium path.

Now, $\max_{x\in S} q^u_{h,L'}(x) - \max_{x\in S} q^u_{h',L'}(x) > 0$, or else concluding search immediately at h would have been a strict improvement. But then by Lemma 12 and induction on the number of searches,

$$\max_{x\in S} q^u_{h,L'}(x) - \max_{x\in S} q^u_{h',L'}(x) < \max_{x\in S} q^u_{h,L''}(x) - \max_{x\in S} q^u_{h',L''}(x).$$

This implies that

$$U(z^*_h, \max_{x\in S} q^u_{h,L'}(x)) - U(z^*_h, \max_{x\in S} q^u_{h',L'}(x)) < U(z^*_h, \max_{x\in S} q^u_{h,L''}(x)) - U(z^*_h, \max_{x\in S} q^u_{h',L''}(x)),$$

by the assumptions that $U_2 \leq 0$ and $U_{22} < 0$. Then in the case where search complexity is L'', no optimal sequential search procedure concludes search at h. This is a contradiction.

Therefore, continuing search is not a part of any optimal sequential search procedure when complexity is L'. □

Proof of Proposition 3. Search outside of the search window is wasteful. Searching a new technology inside the search window but outside of the most promising area results in a quality less than or equal to $z^*_{h_t}$ in the worst case. However, at this history, h_{t+1}, $\max_{x\in S} q^u_{h_t}(x) = \max_{x\in S} q^u_{h_t}(x)$. Therefore, concluding search at h_t would be an improvement for the searcher over reaching h_{t+1} and concluding search. Therefore, if a myopic strategy that continues search at h_t searches insider the most promising area. □

B.4 Proofs from Section 5

Proof of Proposition 4. We prove this result by constructing a candidate z'. Let $q^l_{h_t}$ be the *lower envelope* of Q_{h_t} and let $z' \equiv q^l_{h_t}(\sigma(h_t))$. Not that by definition, $Q_{h'_{t+1}}$ is nonempty, and $z_{t+1} \geq z'$. By construction, and the assumption that $\min_{q\in Q_{h_t}, y\in S} q(y) > 0$, $q^u_{h_{t+1}}(y) = L|\sigma(h_t) - y| + z' \leq q^u_{h_t}(y)$ for all $y \in [x-d, x+d]$, where $d = |x - \sigma(h_t)|$. By Lemma 1, any any search in $[x-d, x+d]$ could not be a part of an optimal search procedure, proving the result. □

Proof of Proposition 5. If the searcher learns good news at $\sigma(h_t)$ and $\sigma(h_t) > x_t$, then $q(x) \leq z_t$ for any $q \in Q_{h_{t+1}}$. Otherwise, q is not quasiconcave. Moreover, $q^u_{h_t}(x) \leq z_t + L(x - x_t)$ for any $x \in S$, by Lemma 2 and the fact that $Q^{QC} \subset Q^{MP}$. Therefore, if $x < x_t + \frac{1}{L}(z_{t+1} - z_t)$, then $q^u_{h_t}(x) < z_{t+1}$. By Lemma 1, any any search in $[\min S, x_t + \frac{1}{L}(z_{t+1} - z_t)]$ cannot not be a part of an optimal search procedure.

The proof of the remaining cases follow identical arguments. □

B.5 Proofs from Section 6

For the minimal prior case (i.e., $\kappa \in (Q^{MP})^k$), it suffices to adapt Lemma 4 to the present case.

Lemma 13. *Consider an attribute learning model where S is a (k-dimensional) rectangular search space, and attributes are substitutes. Let $Q = (Q^{MP})^k$. Fix some history $h'_t = \{(x_i, z'_i)\}_{i=0}^{t-1} \in H$, where $z'_i = (z^{1\prime}_i, \ldots, z^{k\prime}_i)$, and let x be an optimal technology at h'_t. Let $T = X^*_{h'_t} \setminus \{x\}$. For each $x_i \in T$, let $\epsilon_i = (\epsilon^1_i, \ldots, \epsilon^k_i) \in \mathbb{R}^k$. Consider an alternate history $h''_t = \{(x_i, z''_i)\}_{i=0}^{t-1} \in H$ where $z''_i = z'_i + \epsilon_i$ for $x_i \in T$ and $z''_i = z'_i$ otherwise. Suppose moreover that $Q_{h'_t}$ and $Q_{h''_t}$ are nonempty . The searcher's payoff to concluding search in h''_t is greater than or equal to the searcher's payoff to concluding search at h'_t.*

Proof. The nontrivial case is when T is nonempty and some $\epsilon_i \geq 0$. Let $\epsilon = (\epsilon^1, \ldots, \epsilon^k) \in \mathbb{R}^k$ denote the attribute-wise maximum of all the ϵ_i, and let h''' be the history $\{(x_i, z'''_i)\}_{i=0}^{t-1}$, where $z''_i = z'_i + \epsilon$. Let $z', z''' \in [0,1]^k$ denote the attribute-wise maximum of all z'_i's and z'''_i's at histories h' and h''', respectively. Since S is a rectangular search space, there exist

$x', x''' \in S$ such that $\kappa(x') = z'$ and $\kappa(x''') = z'''$. Clearly, x' and x''' are optimal at histories h' and h''', respectively (and the searcher, by assumption, can conclude search with x' (or x'''), even if they were previously unexplored).

Recall, $q(x) = f \circ \kappa(x)$ where f is submodular. Let x'_A denote an adversary optimal unexplored alternative in S, with $z'_A = \kappa(x'_A)$, should the searcher conclude at history h'. Define z'''_A analogously. Now $z'''_A \leq z'_A + \epsilon$ (with equality if the upper bound on quality is not binding).

Now $f(z' + \epsilon) - f(z') > 0$, since f is increasing. Next $z' \leq z'_A$. By submodularity, $f(z' + \epsilon) - f(z') \geq f(z'_A + \epsilon) - f(z'_A)$. In other words, $f(z''') - f(z') \geq f(z'''_A) - f(z'_A)$. The left side of this inequality is the increase in the quality of searcher's best discovery, should he conclude search at history h''' instead of h'. The right side is the increase in quality of the best unexplored option, in the worst case. Along with the assumption that $U_1 \geq -U_2$, this implies that the searcher is better off at history h''' than at h'.

The searcher is weakly better off concluding search at history h'' than at h''', by the same reasoning as in Lemma 4. $\square$

Proof of Corollary 1. Applying Lemma 13 in place of Lemma 4, all other parts of the proof are identical to the proof of Theorem 1. The general case follows from analogues of the proof of Theorem 2. $\square$